arXiv:1510.06463v1 [stat.ML] 22 Oct 2015

Inventory Control Involving Unknown Demand of Discrete Nonperishable Items - Analysis of a Newsvendor-based Policy

Michael N. Katehakis, Jian Yang, and Tingting Zhou
Department of Management Science and information Systems
Business School, Rutgers University, Newark, NJ 07102

October 2015

Abstract

Inventory control with unknown demand distribution is considered, with emphasis placed on the case involving discrete nonperishable items. We focus on an adaptive policy which in every period uses, as much as possible, the optimal newsvendor ordering quantity for the empirical distribution learned up to that period. The policy is assessed using the regret criterion, which measures the price paid for ambiguity on demand distribution over T periods. When there are guarantees on the latter's separation from the critical newsvendor parameter $\beta = b/(h+b)$, a constant upper bound on regret can be found. Without any prior information on the demand distribution, we show that the regret does not grow faster than the rate $T^{1/2+\epsilon}$ for any $\epsilon > 0$. In view of a known lower bound, this is almost the best one could hope for. Simulation studies involving this along with other policies are also conducted.

Keywords: Inventory Control; Newsvendor; Critical Quantile; Empirical Distribution; Large Deviation; Information Theory

1 Introduction

For a given firm, inventory control is about dynamically adjusting ordering quantities to minimize the total long-run expected cost. In traditional models, demand levels faced by a firm are often assumed to be random, however, with known probabilistic distributions. Even the knowledge on demand distribution can often prove to be too optimistic. When the firm has just introduced a new product or when its external environment has just transitioned to a previously unfamiliar phase (such as a severe economic downturn), it will not be sure of the demand patterns to be encountered. One way out is through the Bayesian approach. In it, the firm possesses a prior distribution on potential demand patterns. Then, posterior understanding on demand is updated by its realized levels. Inventory management taking this approach can be found in, for instance, Scarf [20] and Lariviere and Porteus [15].

Most other times, even a prior distribution on demand can seem far-fetched. What meager information one possesses might just be a collection of potential demand distributions. Now, the concerned firm has still to make decisions based on its past observations. But its goal is no longer about catering to a specific demand distribution or even a series of posterior demand distributions. Rather, its history-dependent (henceforward called adaptive) control policy should better yield results that are reasonably good under all potential demand distributions from the given collection. A given policy's regret under a given demand distribution and over a fixed time horizon measures the price paid for ambiguity; namely, it registers the difference between the policy's performance and that of the best policy tailor-made for the demand distribution had it been known. A policy will be considered good when its worst regret over all demand distributions in a collection grows as slowly as possible over time.

In this paper we follow the frequentist approach. This approach was pioneered in the work of Robbins [18], Lai and Robbins [14], Katehakis and Robbins [12], and Burnetas and Katehakis [3], for allocation problems, where the main concern is on adaptively selecting the most promising population to draw from so as to maximize the total expected value of samples, or equivalently to minimize the regret due to ignorance of the distributions. Later, the approach was brought to bear on adaptive Markov decision processes; see, e.g., Burnetas and Katehakis [4]. For recent work in this area we refer to Burnetas et al [5], and Cowan and Katehakis [6], [7].

Adaptive policies for inventory control have been considered. Huh and Rusmevichientong [11] analyzed a gradient-based policy most suitable to the continuous-demand case. Their policy could also be thought of as an extension of stochastic approximation (SA), which

started with Robbins and Monro [19] and Kiefer and Wolfowitz [13]. More recently, Besbes and Muharremoglu [2] focused on the discrete-demand case of the repeated newsvendor problem and proposed policies with provably good performance guarantees.

In a revenue management setup, Besbes and Zeevi [1] studied the dynamic setting of prices while learning demand on the fly. Perakis and Roels [17] minimized the worst regret for a single-period newsvendor problem. Rather than dynamic learning, they used conic optimization to deal with partial information on random demands in the form of known moments. For the newsvendor problem and its multi-period version involving nonperishable items, Levi, Roundy, and Shmoys [16] relied on randomly generated demand samples to reach solutions with relatively good qualities at high probabilities. In this work, demand learning was through a black box capable of turning up an arbitrary number of samples at any time, rather than through the sequence of actually encountered demand levels.

We study inventory control involving the online learning of unknown demand, focusing on nonperishable items like Huh and Rusmevichientong [11] and discrete-quantity analysis like Besbes and Muharremoglu [2]. In many real-life situations ranging from manufacturing to retailing, nonperishability of items is an essential feature to be faced head on. Also, many applications, such as the management of bulk items, dictate that demand be discrete. An adaptation to Huh and Rusmevichientong's [11] policy could work for the discrete-demand case, as demonstrated in the repeated-newsvendor analysis in their Section 3.4. However, further adaption seems needed for the nonperishable-inventory case; see (45) and (46) later in our simulation study. To the best of our knowledge, our theoretical performance guarantees on an adaptive inventory policy involving unknown demand of discrete nonperishable items have made original contributions.

We adopt a very simple and natural policy, the one that always orders, as much as possible, to the critical newsvendor quantile corresponding to the empirical demand distribution. The optimal ordering quantity for a newsvendor problem involving holding cost h, backlogging or effective lost sales cost b, and known demand distribution f, is the β-quantile of the distribution f, where $\beta = b/(h+b)$. By the beginning of period t, the empirical distribution $\hat{f}_{t-1}$ one has about demand is defined by the frequencies of various demand levels reached in periods $1, 2, ..., t-1$. The heuristic policy advocates ordering up to the β-quantile of $\hat{f}_{t-1}$ in every period t. It has been considered by Besbes and Muharremoglu [2]. Our nonperishable-inventory variant needs to further ensure that items left over from earlier periods are accounted for. This is a difficult point that requires quite careful treatments.

We analyze two cases. In the first case in which the β-quantile estimation of none of the demand distributions considered is overly sensitive to small errors, we show that the worst regret will be bounded by a time-invariant constant. Though the flat rate over time is impressive, the result can be faulted by its requirement on the unknown distribution's behavior around its β-quantile. We thus go on to the second, more involved, case where all potential demand distributions are allowed. Given any $\epsilon > 0$, we show that the worst regret over all distributions will not grow faster than the rate $T^{1/2+\epsilon}$. In view of the $T^{1/2}$-sized lower bound achieved at Lemma 4 of Besbes and Muharremoglu [2], this is almost the best one could hope for. Our derivation invokes large-deviation and information-theoretic results such as Sanov's Theorem and Pinsker's inequality, as well as innate features of empirical distributions and inventory control. Methodological advances might find applications elsewhere.

Our simulation study indicates the high likelihood with which the worst regret grows at the $T^{1/2}$-rate. Thus, it remains as a future research item whether the $\epsilon > 0$ in our upper bound can be removed. The study also shows that the newsvendor-based policy compares favorably with the adaptation to Huh and Rusmevichientong's [11] SA-based policy. This is expected, as the former uses more information about past demand realizations. Through the study, we also confirm that the underlying distribution's separation from β plays a prominent role in determining the regret generated by the newsvendor-based policy.

The remainder of the paper is organized as follows. Section 2 introduces notation, problem formulation, and the newsvendor-based policy. The time-invariant bound with demand restriction around β is given in Section 3, while the slower-over-time bound without any demand restriction is derived in Section 4. We present results of our simulation study in Section 5. The paper is concluded in Section 6.

2 Problem and Policy

We consider a multi-period inventory control problem in which unsatisfied demand is either backlogged or lost. Also, items are nonperishable so that those unsold in one period are carried over to the next period. Demand D_t in each period $t = 1, 2, ...$ is a random draw from a distribution with discrete support $\{0, 1, ..., \bar{d}\}$, where $\bar{d}$ is some positive integer. A generic distribution is representable by a vector $f \equiv (f(0), f(1), ..., f(\bar{d}))$ in the $\bar{d}$-dimensional

simplex within $\mathbb{R}^{\bar{d}+1}$:

$$\Delta \equiv \{f \in [0,1]^{\bar{d}+1} | \sum_{d=0}^{\bar{d}} f(d) = 1\}. \tag{1}$$

We use F_f to denote the cumulative distribution function (cdf) associated with any given $f \in \Delta$. It satisfies $F_f(d) = \sum_{d'=0}^{d} f(d')$ for $d' = 0, 1, ..., \bar{d}$.

We suppose production cost is linear at a unit rate c. The concerned planning horizon constitutes periods $1, 2, ..., T$. In the terminal period T, the firm will gain cx if it has $x = 1, 2, ...$ leftover items and will need to pay $-cx$ if $x = -1, -2, ...$ Also, assume strictly positive holding cost rate h and strictly positive backlogging or effective lost sales cost rate b. In the backlogging case, b is usually in the same order of magnitude as h; whereas, in the lost sales case, b is in the order of an item's profit margin and normally significantly greater than h. When the firm orders up to y_t and experiences demand d_t in each period $t = 1, 2, ..., T$, its total cost in the backlogging case will be

$$c \cdot \sum_{t=1}^{T} d_t + \sum_{t=1}^{T} [h \cdot (y_t - d_t)^+ + b \cdot (d_t - y_t)^+]. \tag{2}$$

Since the first term in the above is not affected by the decision sequence $(y_1, y_2, ..., y_T)$, we shall focus on the latter inventory-related cost term. In the lost sales case, the same conclusion can be reached when b is an item's profit margin plus the actual per-item lost sales cost; see Huh and Rusmevichientong [11].

Given $f \in \Delta$, $t = 1, 2, ...$, and real-valued function g defined on $\{0, 1, ..., \bar{d}\}^t$, we use $\mathbb{E}_f[g(D_1, ..., D_t)]$ to represent the average of $g(d_1, ..., d_t)$ when each D_s is independently sampled from distribution f:

$$\mathbb{E}_f[g(D_1, ..., D_t)] \equiv \sum_{d_1=0}^{\bar{d}} \cdots \sum_{d_t=0}^{\bar{d}} f(d_1) \times \cdots \times f(d_t) \times g(d_1, ..., d_t). \tag{3}$$

For subset $D' \subset \{0, 1, ..., \bar{d}\}^t$, we understand $\mathbb{P}_f[D']$ by $\mathbb{E}_f[\mathbf{1}((D_1, ..., D_t) \in D')]$. Note that

$$\mathbb{P}_f[D_t \leq d] = \mathbb{E}_f[\mathbf{1}(D_t \leq d)] = \sum_{d'=0}^{d} f(d') = F_f(d). \tag{4}$$

Define $Q_f(y)$ for every $f \in \Delta$ and $y = 0, 1, ..., \bar{d}$, so that

$$\begin{aligned} Q_f(y) &\equiv \mathbb{E}_f[h \cdot (y - D)^+ + b \cdot (D - y)^+] = \textstyle\sum_{d=0}^{\bar{d}} f(d) \cdot (h \cdot (y - d)^+ + b \cdot (d - y)^+) \\ &= h \cdot \textstyle\sum_{d=0}^{y-1} F_f(d) + b \cdot \sum_{d=y}^{\bar{d}-1} (1 - F_f(d)). \end{aligned} \tag{5}$$

It is the single-period average cost under order-up-to level y. Let $Q_f^* = \min_{y=0,1,...,\bar{d}} Q_f(y)$ be the minimum cost in one period under f. Suppose y_f^* is an order-up-to level that achieves the one-period minimum. Then, when facing a T-period horizon, an optimal policy will be to repeatedly order up to this level. Thus, the minimum cost over T periods is $Q_f^* \cdot T$.

A salient feature of our current problem, however, is that f is not known beforehand. So instead of any f-dependent policy, we seek a good f-independent policy which takes advantage of demand levels observed in the past. An adaptive policy $\mathbf{y} \equiv (y_1, y_2, ...)$ is such that, for $t = 1, 2, ...$, each $y_t = 0, 1, ..., \bar{d}$ is a function of the historical demand vector $\mathbf{d}_{[1,t-1]} \equiv (d_1, ..., d_{t-1}) \in \{0, 1, ..., \bar{d}\}^{t-1}$. Under it, the T-period total average cost is

$$Q_f^T(\mathbf{y}) \equiv \sum_{t=1}^{T} \mathbb{E}_f[h \cdot (y_t(\mathbf{D}_{[1,t-1]}) - D_t)^+ + b \cdot (D_t - y_t(\mathbf{D}_{[1,t-1]}))^+] = \sum_{t=1}^{T} \mathbb{E}_f[Q_f(y_t(\mathbf{D}_{[1,t-1]}))], \quad (6)$$

where the second equality comes from the independence between $\mathbf{D}_{[1,t-1]}$ and D_t. Now define T-period regret $R_f^T(\mathbf{y})$ of using the adaptive policy $\mathbf{y}$ against the unknown distribution f:

$$R_f^T(\mathbf{y}) \equiv Q_f^T(\mathbf{y}) - Q_f^* \cdot T = \sum_{t=1}^{T} \mathbb{E}_f[Q_f(y_t(\mathbf{D}_{[1,t-1]}))] - Q_f^* \cdot T. \quad (7)$$

Here, the ultimate goal should be that of identifying adaptive policies $\mathbf{y}$ that prevent $R_f^T(\mathbf{y})$ from growing too fast in T under all or at least most f's within Δ.

We concentrate on one policy inspired by an optimal y_f^* when f is known. From (5), we see that necessary and also sufficient conditions for optimality of any y are

$$Q_f(y+1) - Q_f(y) = (h+b) \cdot F_f(y) - b \geq 0, \quad (8)$$

and

$$Q_f(y) - Q_f(y-1) = (h+b) \cdot F_f(y-1) - b \leq 0, \quad (9)$$

Let $\beta = b/(h+b)$ be the famous newsvendor parameter that lies in $(0, 1)$. For $f \in \Delta$, let $y_{f,\beta}^*$ be the associated newsvendor order-up-to level, so that

$$y_{f,\beta}^* \equiv F_f^{-1}(\beta) \equiv \min\{d = 0, 1, ..., \bar{d} | F_f(d) \geq \beta\}. \quad (10)$$

By definition, $F_f(y_{f,\beta}^*) \geq \beta$ and hence $Q_f(y_{f,\beta}^*+1) - Q_f(y_{f,\beta}^*) \geq 0$ by (8); also, $F_f(y_{f,\beta}^*-1) < \beta$ and hence $Q_f(y_{f,\beta}^*) - Q_f(y_{f,\beta}^* - 1) < 0$ by (9). Therefore, $Q_f(y_{f,\beta}^*) = Q_f^*$, meaning that $y_{f,\beta}^*$ is an optimal order-up-to level for the one-period problem when f is known. Now with f unknown, we might adopt level $y_{f_{t-1},\beta}^*$ where f_{t-1} is a good estimate of f after observing demand vector $D_{[1,t-1]}$. .

The prime candidate for f_{t-1} is the empirical distribution $\hat{f}_{t-1}$. For $t = 2, 3, ...$, define $\hat{f}_{t-1} \in \Delta$ by $(\hat{f}_{t-1}(0), \hat{f}_{t-1}(1), ..., \hat{f}_{t-1}(\bar{d}))$, so that for every $d = 0, 1, ..., \bar{d}$,

$$\hat{f}_{t-1}(d) = \frac{\sum_{s=1}^{t-1} \mathbf{1}(d_s = d)}{t-1}. \tag{11}$$

Each $\hat{f}_{t-1}$ has its corresponding cdf $\hat{F}_{t-1} \equiv F_{\hat{f}_{t-1}}$. Both $\hat{f}_{t-1}$ and $\hat{F}_{t-1}$ are certainly functions of the past demand vector $d_{[1,t-1]} \equiv (d_1, d_2, ..., d_{t-1})$. However, for notational simplicity we have refrained from making this dependence explicit. Our heuristic policy applies the newsvendor formula to the empirical demand distribution. It lets the firm order nothing in period 1; that is, $y_1 = \hat{y}_1 = 0$. For any $t = 2, 3, ...$, it advises the firm to order up to

$$y_t = \hat{y}_t \vee (y_{t-1} - d_{t-1}), \tag{12}$$

in period t, in which

$$\hat{y}_t = y^*_{\hat{f}_{t-1},\beta} = \hat{F}^{-1}_{t-1}(\beta) \equiv F^{-1}_{\hat{f}_{t-1}}(\beta) = \min\{d = 0, 1, ..., \bar{d} | \sum_{s=1}^{t-1} \mathbf{1}(d_s \geq d) \geq \beta \cdot (t-1)\}. \tag{13}$$

For simplicity, we have not made the dependence of y_t and $\hat{y}_t$ on $d_{[1,t-1]}$ explicit. Nor have we used the full $d_{[1,t-2]}$-dependent notation on y_{t-1}. This will apply to the remainder of the paper. For the lost sales case, we need to guarantee that $y_t \geq 0$ and hence enhance (12) to $y_t = \hat{y}_t \vee (y_{t-1} - d_{t-1})^+$. However, the current heuristic through (13) has ensured that $\hat{y}_t \geq 0$. So the same (12) can still be used. Of course, a typical β for the backlogging case might be around 1/2, whereas a typical β for the lost sales case might be close to 1.

Our main purpose is to show that the f-blind and yet adaptive policy $\mathbf{y}$ described by (12) and (13) will incur regret $R^T_f(\mathbf{y})$ as defined by (7) that is slow-growing in the planning length T for most or even all f's among Δ. The requirement of $y_t \geq y_{t-1} - d_{t-1}$ in (12), as necessitated by our nonperishable-inventory setting, renders decisions made over different periods more entangled with one another. Along with the discrete-demand setup, this substantially complicates the problem's analysis.

3 Separation-affected Bound

We first establish a bound for $R^T_f(\mathbf{y})$ when there are guarantees on the distances between the $F_f(d)$ values and the critical value β. By (7) and (12),

$$R^T_f(\mathbf{y}) = R^{T1}_f(\mathbf{y}) + R^{T2}_f(\mathbf{y}), \tag{14}$$

where

$$R_f^{T1}(\mathbf{y}) = \sum_{t=1}^{T} \mathbb{E}_f[Q_f(\hat{y}_t)] - Q_f^* \cdot T, \tag{15}$$

and, since $y_1 = \hat{y}_1 = 0$ by design and hence $y_2 = \hat{y}_2$,

$$R_f^{T2}(\mathbf{y}) = \sum_{t=3}^{T} \mathbb{E}_f[Q_f(y_t) - Q_f(\hat{y}_t)]. \tag{16}$$

It might be said that $R_f^{T1}(\mathbf{y})$ represents the price paid for the regrettable fact that the policy $\mathbf{y}$ was not designed with the particular distribution f in mind; meanwhile, $R_f^{T2}(\mathbf{y})$ captures the additional cost due to the nonperishability of items. We find it convenient to bound $R_f^{T1}(\mathbf{y})$ and $R_f^{T2}(\mathbf{y})$ separately.

Due to $y_{f,\beta}^*$'s definition in (10), we must have $f(y_{f,\beta}^*) > 0$; for otherwise, $y_{f,\beta}^*$ could be made even smaller. Now there are two cases, with

Case 1: $\alpha_{f,\beta} \equiv F_f(y_{f,\beta}^* - 1) < \beta < \gamma_{f,\beta} \equiv F_f(y_{f,\beta}^*)$; and,

Case 2: $\alpha_{f,\beta} \equiv F_f(y_{f,\beta}^* - 1) < \beta = F_f(y_{f,\beta}^*)$.

We concentrate on case 1 first. Let us use $x \in [0,1]$ to denote the Bernoulli distribution where the chance for 1 is x and that for 0 is $1-x$. Then, the numbers $\alpha_{f,\beta} = F_f(y_{f,\beta}^* - 1)$ and $\gamma_{f,\beta} = F_f(y_{f,\beta}^*)$ represent two Bernoulli distributions. Use $\hat{\alpha}_{t-1}$ and $\hat{\gamma}_{t-1}$ for the empirical distributions of $\alpha_{f,\beta}$ and $\gamma_{f,\beta}$, respectively, both recording frequencies of 1's in the first $t-1$ Bernoulli draws. One important observation is that

$$\begin{cases} \mathbb{P}_f[\hat{F}_{t-1}^{-1}(\beta) \le y_{f,\beta}^* - 1] = \mathbb{P}_f[\hat{F}_{t-1}(y_{f,\beta}^* - 1) \ge \beta] = \mathbb{P}_{\alpha_{f,\beta}}[\hat{\alpha}_{t-1} \ge \beta], \\ \mathbb{P}_f[\hat{F}_{t-1}^{-1}(\beta) \ge y_{f,\beta}^* + 1] = \mathbb{P}_f[\hat{F}_{t-1}(y_{f,\beta}^*) < \beta] = \mathbb{P}_{\gamma_{f,\beta}}[\hat{\gamma}_{t-1} < \beta]. \end{cases} \tag{17}$$

By a special version of Sanov's Theorem (Dembo and Zeitouni [9], (2.1.12)), we have upper bounds for right-hand sides above:

$$\begin{cases} \mathbb{P}_{\alpha_{f,\beta}}[\hat{\alpha}_{t-1} \ge \beta] \le t^2 \cdot \exp(-(t-1) \cdot \inf_{x \in [\beta,1]} \mathrm{D}_{KL}(x||\alpha_{f,\beta})), \\ \mathbb{P}_{\gamma_{f,\beta}}[\hat{\gamma}_{t-1} < \beta] \le t^2 \cdot \exp(-(t-1) \cdot \inf_{x \in [0,\beta)} \mathrm{D}_{KL}(x||\gamma_{f,\beta})), \end{cases} \tag{18}$$

respectively, where

$$\mathrm{D}_{KL}(u||v) \equiv u \cdot \ln(\frac{u}{v}) + (1-u) \cdot \ln(\frac{1-u}{1-v}), \tag{19}$$

i.e., the relative entropy or Kullback-Leibler divergence between Bernoulli distributions u and v. Since $\mathrm{D}_{KL}(\cdot||v)$ is known to be convex with minimum achieved at v (Cover and Thomas [8], Theorems 2.6.3 and 2.7.2),

$$\begin{cases} \inf_{x \in [\beta,1]} \mathrm{D}_{KL}(x||\alpha_{f,\beta}) = \mathrm{D}_{KL}(\beta||\alpha_{f,\beta}) > 0, \\ \inf_{x \in [0,\beta)} \mathrm{D}_{KL}(x||\gamma_{f,\beta}) = \mathrm{D}_{KL}(\beta||\gamma_{f,\beta}) > 0. \end{cases} \tag{20}$$

Now

$$\kappa_{f,\beta} \equiv \min\{\mathrm{D}_{KL}(\beta||\alpha_{f,\beta}), \mathrm{D}_{KL}(\beta||\gamma_{f,\beta})\} > 0. \tag{21}$$

Combining (17) to (21), we have both $\mathbb{P}_f[\hat{F}_{t-1}^{-1}(\beta) \leq y_{f,\beta}^* - 1]$ and $\mathbb{P}_f[\hat{F}_{t-1}^{-1}(\beta) \geq y_{f,\beta}^* + 1]$ being bounded from above by $t^2 \cdot \exp(-\kappa_{f,\beta} \cdot (t-1))$. In view of (13),

$$\mathbb{P}_f[\hat{y}_t \leq y_{f,\beta}^* - 1] \vee \mathbb{P}_f[\hat{y}_t \geq y_{f,\beta}^* + 1] \leq \min\{t^2 \cdot \exp(-\kappa_{f,\beta} \cdot (t-1)), 1\}. \tag{22}$$

We can identify a large enough $\tau_{f,\beta}$, so that for $t \geq \tau_{f,\beta} + 1$, both

$$t^2 \cdot \exp(-\kappa_{f,\beta} \cdot (t-1)) < \frac{1}{2}, \tag{23}$$

and

$$\frac{(t+1)^2 \cdot \exp(-\kappa_{f,\beta} \cdot t)}{t^2 \cdot \exp(-\kappa_{f,\beta} \cdot (t-1))} < \exp(-\frac{\kappa_{f,\beta}}{2}). \tag{24}$$

Now $R_f^{T1}(\mathbf{y})$ as given in (15) can be bounded.

Proposition 1 *Under case 1, it is true that*

$$R_f^{T1}(\mathbf{y}) \leq (h\bar{d} + b\bar{d}) \cdot (\tau_{f,\beta} + \frac{1}{2 \cdot (1 - \exp(-\kappa_{f,\beta}/2))}).$$

All proofs of this section can be found in Appendix A. Proposition 1 demonstrates that there is a T-independent bound for $R_f^{T1}(\mathbf{y})$; however, the constant is heavily dependent on the relative positioning between the unknown distribution f and the parameter β. It conveys the same message as Besbes and Muharremoglu's [2] Theorem 2 on the repeated newsvendor problem. For our case involving nonperishable items, we still need a bound on $R_f^{T2}(\mathbf{y})$ defined in (16). For this purpose, suppose $\epsilon_f \equiv f(\bar{d}) > 0$.

Proposition 2 *Under case 1, it is true that*

$$R_f^{T2}(\mathbf{y}) \leq h\bar{d} \cdot (\tau_{f,\beta} + \frac{1}{1 - \exp(-\kappa_{f,\beta}/2)} + \frac{1 - \epsilon_f}{\epsilon_f} + \frac{1}{2\epsilon_f \cdot (1 - \exp(-\kappa_{f,\beta}/2))}).$$

The proof revolves around bounding $\mathbb{P}_f[y_{t+1} \geq \hat{y}_{t+1} + 1]$ for t large enough. When case 2 occurs, let $y_{f,\beta}^1 \geq y_{f,\beta}^*$ be such that

$$\beta = F_f(y_{f,\beta}^*) = F_f(y_{f,\beta}^* + 1) = \cdots = F_f(y_{f,\beta}^1) < F_f(y_{f,\beta}^1 + 1). \tag{25}$$

From (8), we see that

$$Q_f(y_{f,\beta}^*) = Q_f(y_{f,\beta}^* + 1) = \cdots = Q_f(y_{f,\beta}^1 + 1) < Q_f(y_{f,\beta}^1 + 2). \tag{26}$$

So instead of $\mathbb{P}_f[\hat{y}_t \geq y^*_{f,\beta} + 1]$, we need only to estimate $\mathbb{P}_f[\hat{y}_t \geq y^1_{f,\beta} + 2]$. But according to (17), this is the same as $\mathbb{P}_{\gamma_{f,\beta}}[\hat{\gamma}_{t-1} < \beta]$ for $\gamma_{f,\beta} = F_f(y^1_{f,\beta} + 1)$. Now due to (25), $\min\{\beta - \alpha_{f,\beta}, \gamma_{f,\beta} - \beta\} = \min\{f(y^*_{f,\beta}), f(y^1_{f,\beta} + 1)\} > 0$. So the same bounds as in Propositions 1 and 2 can be established. A definition suitable for both cases is that

$$\alpha_{f,\beta} \equiv \max\{F_f(d)|F_f(d) < \beta\}, \qquad \gamma_{f,\beta} \equiv \min\{F_f(d)|F_f(d) > \beta\}. \tag{27}$$

Now we can achieve an upper bound for $R^T_f(\mathbf{y})$, regardless of the case we are in. For any $f \in \Delta$, define f's separation from β by

$$\delta_{f,\beta} \equiv (\beta - \alpha_{f,\beta}) \wedge (\gamma_{f,\beta} - \beta). \tag{28}$$

Given $\epsilon, \delta > 0$, define $\Delta_{\epsilon,\delta,\beta} \subset \Delta$ to be the collection of f's with guaranteed lower bounds on both ϵ_f and $\delta_{f,\beta}$:

$$\Delta_{\epsilon,\delta,\beta} \equiv \{f \in \Delta | \epsilon_f \geq \epsilon, \delta_{f,\beta} \geq \delta\}. \tag{29}$$

It turns out that $R^T_f(\mathbf{y})$ has an upper bound that is uniform across $f \in \Delta_{\epsilon,\delta,\beta}$.

Theorem 1 *For any $\epsilon, \delta > 0$, there is a positive constant $A_{\epsilon,\delta,\beta}$ so that*

$$\sup_{f \in \Delta_{\epsilon,\delta,\beta}} R^T_f(\mathbf{y}) \leq A_{\epsilon,\delta,\beta}.$$

Theorem 1 shows that the regret $R^T_f(\mathbf{y})$ is bounded from above by a constant independent of time T, so long as there are known lower bounds on both $\epsilon_f \equiv f(\bar{d})$ and $\delta_{f,\beta}$, the separation between f and β. Between the two requirements, the first one appears more reasonable as $\bar{d}$ can always be the highest level that demand can ever reach. The second requirement, on the other hand, straddles between both demand distributions and cost parameters. It seems far-fetched to exclude a priori distributions f satisfying $\delta_{f,\beta} \in (0, \delta)$ from consideration.

4 Universal Bound

Due to the shortcoming inherent in the previous bound, we feel compelled to derive a bound on $R^T_f(\mathbf{y})$ that requires no prior knowledge on f, let alone its relative positioning with respect to the cost parameter β. The new analysis involves Sanov's Theorem which offers a large-deviational bound on the difference between an empirical distribution and its generating distribution, Pinsker's inequality which connects two distances between distributions, Markov's inequality, and innate properties of the inventory management problem.

Let $\mathrm{D}_{KL}(g||f)$ be the relative entropy or Kullback-Leibler divergence between distributions g and f in Δ, so that

$$\mathrm{D}_{KL}(g||f) \equiv \sum_{d=0}^{\bar{d}} g(d) \cdot \ln(\frac{g(d)}{f(d)}). \tag{30}$$

Sanov's Theorem (Dembo and Zeitouni [9], (2.1.12)) states that, for any set G of the demand space Δ that is closed in the Euclidean metric,

$$\mathbb{P}_f[\hat{f}_{t-1} \in G] \leq t^{\bar{d}+1} \cdot \exp(-(t-1) \cdot \inf_{g \in G} \mathrm{D}_{KL}(g||f)). \tag{31}$$

Let $\delta_V(f,g)$ be the total variation between distributions f and g; i.e.,

$$\delta_V(f,g) \equiv \max_{K \subset \{0,1,\ldots,\bar{d}\}} \mid \sum_{d \in K} f(d) - \sum_{d \in K} g(d) \mid, \tag{32}$$

which also equals $|| f - g ||_1 /2 \equiv \sum_{d=0}^{\bar{d}} \mid f(d) - g(d) \mid /2$. Pinsker's inequality (Cover and Thomas [8], Lemma 11.6.1) specifies that

$$\delta_V(f,g) \leq \sqrt{\frac{\mathrm{D}_{KL}(g||f)}{2}}. \tag{33}$$

With the absence of any prior knowledge on the $f \in \Delta$, we can manage to obtain a $T^{1/2+\epsilon}$-bound on the $R_f^{T1}(\mathbf{y})$ defined in (15) for any $\epsilon > 0$. Due to (13), the key is to show that $Q_f(y^*_{\hat{f}_{t-1},\beta}) - Q_f(y^*_{f,\beta})$ will converge to 0 quickly. In view of the Sanov property (31), this will be achievable if we can show that $Q_f(y^*_{g,\beta}) - Q_f(y^*_{f,\beta})$ will be small when g and f are close by. This is when Pinsker's inequality (33) and other properties related to the inventory management problem, such as the optimality of $y^*_{f,\beta}$ to $Q_f(\cdot)$ and the linearity of $Q_f(y) - Q_g(y)$ in the distance between f and g, will be useful. The final form of the bound comes from the estimation of certain summations through integrations.

Proposition 3 *For any $\epsilon > 0$, there are positive constants A_ϵ and B_ϵ, so that*

$$R_f^{T1}(\mathbf{y}) \leq A_\epsilon + B_\epsilon \cdot T^{1/2+\epsilon}.$$

All proofs of this section can be found in Appendix B. The current result is ultimately about bounding the two sums $\sum_{t=1}^{T} \sqrt{\varepsilon_t}$ and $\sum_{t=1}^{T} t^{\bar{d}+1} \cdot \exp(-\varepsilon_t \cdot (t-1))$ simultaneously for a sequence ε_t; see (B.8) in our proof. The $T^{1/2+\epsilon}$-sized rate comes from our choice of $\varepsilon_t = (t-1)^{-1+2\epsilon}$ to balance the two terms. If we remove $\epsilon > 0$, the first sum will give a

$T^{1/2}$-sized rate but the second sum will fail to achieve a below-T rate. This is why we can attain the $T^{1/2+\epsilon}$-sized rate for an arbitrarily small $\epsilon > 0$ but never the exact $T^{1/2}$-sized rate.

Next, we obtain a bound in the same order of magnitude for $R_f^{T2}(\mathbf{y})$ as defined by (16). Let us first get some more understanding on the entity. From (12), we see that

$$y_t = \hat{y}_t \vee (\hat{y}_{t-1} - d_{t-1}) \vee (\hat{y}_{t-2} - d_{t-2} - d_{t-1}) \vee \cdots \vee (\hat{y}_1 - d_1 - d_2 - \cdots - d_{t-1}). \tag{34}$$

There is a latest s so that

$$y_t = \hat{y}_s - d_s - d_{s+1} - \cdots - d_{t-1}, \tag{35}$$

which occurs exactly when

$$\hat{y}_s - d_s - d_{s+1} - \cdots - d_{t-1} - 1 \geq \hat{y}_t \vee (\hat{y}_{t-1} - d_{t-1}) \vee \cdots \vee (\hat{y}_{s+1} - d_{s+1} - \cdots - d_{t-1}), \tag{36}$$

and

$$\hat{y}_s \geq (\hat{y}_{s-1} - d_{s-1}) \vee (\hat{y}_{s-2} - d_{s-2} - d_{s-1}) \vee \cdots \vee (\hat{y}_1 - d_1 - d_2 - \cdots - d_{s-1}). \tag{37}$$

Inspired by the above, we define random variables $I \geq 1$ and $S_1, S_2, ..., S_I, S_{I+1}$ in an iterative fashion as follows. First, let $S_1 = 1$. Now for some $i = 1, 2, ...$, suppose S_i has been settled. Then, let S_{i+1} be the first t after S_i so that

$$\hat{y}_t \geq \hat{y}_{S_i} - D_{S_i} - D_{S_i+1} - \cdots - D_{t-1}, \tag{38}$$

if such a $t \leq T$ can be identified. If not, mark the latest i as I and let $S_{I+1} = T + 1$. For any t, let $L(t)$ be the largest $S_i \leq t$. This $L(t)$ can serve as the earlier s satisfying (36) and (37) that corresponds to t. Note that $L(t)$ along with $D_{L(t)}, D_{L(t)+1}, ..., D_{t-1}$ are independent of D_t. So by (16), as well as (35) to (38),

$$R_f^{T2}(\mathbf{y}) = \sum_{t=3}^{T} \mathbb{E}_f[Q_f(\hat{y}_{L(t)} - D_{L(t)} - \cdots - D_{t-1}) - Q_f(\hat{y}_t)]. \tag{39}$$

Now we are in a position to derive the bound.

Proposition 4 *For any $\epsilon > 0$, there exist positive constants C_ϵ and D_ϵ so that*

$$R_f^{T2}(\mathbf{y}) \leq C_\epsilon + D_\epsilon \cdot T^{1/2+\epsilon}.$$

This is the most difficult result of the paper. We will exploit the observations made from (34) to (39) to the fullest extent, with the basic understanding that the actual order-up-to level y_t will be $\hat{y}_s - D_s - \cdots - D_{t-1}$ for some $s \leq t$. We are tasked to show that the term $Q_f(\hat{y}_s - D_s - \cdots - D_{t-1}) - Q_f(\hat{y}_t)$ can be bounded. For $\gamma = 1 - f(0)$, we divide the proof into two cases, the one with $\gamma \geq (1-\beta)/2$ and the other one with $\gamma < (1-\beta)/2$.

In the former large-γ case, demand will accumulate over time with a guaranteed speed and $\hat{y}_t \geq \hat{y}_s - D_s - \cdots - D_{t-1}$ will occur ever more surely as $t - s$ increases. Then, for the minority case where $t - s$ is small, by exploiting natures of the empirical distribution and the newsvendor formula, we can come up with bounds related to $|Q_f(\hat{y}_s - D_s - \cdots - D_{t-1}) - Q_f(\hat{y}_t)| \cdot \mathbf{1}(\hat{y}_t \leq \hat{y}_s - D_s - \cdots - D_{t-1} - 1)$. Especially important is the observation that $\hat{y}_t \leq \hat{y}_s - D_s - \cdots - D_{t-1} - 1$ only if

$$\beta \leq \hat{F}_{t-1}(\hat{y}_t) \leq \hat{F}_{t-1}(\hat{y}_s - D_s - \cdots - D_{t-1} - 1) < \beta + \frac{t-s}{s}; \tag{40}$$

see (B.25) later. We will end up with a trade-off already encountered in the proof of Proposition 3; see (B.31) in the proof. This is the source of the $T^{1/2+\epsilon}$-sized growth rate. However, the constants will grow as γ shrinks, because it takes ever longer for demand to accumulate.

Therefore, we seek a different approach for the latter small-γ case, when $\gamma < (1-\beta)/2 < 1-\beta$. This is the time when $y^*_{f,\beta} = \hat{F}_f^{-1}(\beta) = 0$ because $F_f(0) = f(0) = 1-\gamma > \beta$. We utilize the fact that $\hat{y}_s - D_s - \cdots - D_{t-1} \geq 1$ is the bare minimum for $\hat{y}_s - D_s - \cdots - D_{t-1} \geq \hat{y}_t + 1$. But the latter will be true only if both $\hat{y}_s \geq 1$ and for some $d = 1, 2, ..., \bar{d}$, both $\hat{y}_s \geq d$ and $D_s + \cdots + D_{t-1} \leq d-1$. For all γ's in the interval $[0, (1-\beta)/2)$, we achieve a uniform bound in the order of $\ln T$, which is dominated by the one obtained in the first case.

Besides Sanov's Theorem and Pinsker's inequality, the proof also exploits Markov's inequality in bounding $\mathbb{P}_f[\hat{y}_s \geq d]$, which, through (10) to (13), is the chance for the portion of earlier demand levels at or exceeding d to be greater than $1 - \beta$. Combining Propositions 3 and 4, we get a bound for $R_f^T(\mathbf{y})$ that is not tangled up with how f positions with β.

Theorem 2 *For any $\epsilon > 0$, there are positive constants E_ϵ and F_ϵ so that*

$$\sup_{f \in \Delta} R_f^T(\mathbf{y}) \leq E_\epsilon + F_\epsilon \cdot T^{1/2+\epsilon}.$$

The constants involved can depend on the problem's parameters h, b, and $\bar{d}$. However, they are uniform across all f's in Δ. For the repeated newsvendor problem, Besbes and Muharremoglu [2] has already shown a $T^{1/2}$-sized lower bound (Lemma 4, with ε replaced

by $1/\sqrt{T}$ in its (C-8)). The example used for the bound involves distributions f with small separations from β but in opposite directions. According to (12), the current case merely adds the restriction $y_t(\mathbf{d}_{[1,t-1]}) \geq y_{t-1}(\mathbf{d}_{[1,t-2]}) - d_{t-1}$ to the adaptive policy considered. So the lower bound can be no better.

In view of this, the above is almost the best one can hope for. Huh and Rusmevichientong's [11] SA-based policy was shown to have a $T^{1/2}$-sized bound when ordering quantities are discrete. But the analysis was done for the repeated-newsvendor setting. The policy's adaptation to the nonperishable-item setting, as to occur in (45) and (46), appears to be more complicated. Its full analysis awaits further research.

5 Simulation Study

In the study, we fix $h+b = 10$ and $\bar{d} = 20$. We let b's variation drive changes in $\beta = b/(h+b)$. Smaller β values are suitable for the backlogging case and larger β values the lost sales case. At each parameter combination, we randomly generate $K = 1,000$ distributions f in Δ uniformly. Within each run $k = 1, 2, ..., K$, particularly, let $\xi_1, ..., \xi_{\bar{d}}$ be independent random variables uniformly distributed on $[0, 1]$. Rank them so that

$$\eta_1 \equiv \xi_{(1)} < \eta_2 \equiv \xi_{(2)} < \cdots < \eta_{\bar{d}} \equiv \xi_{(\bar{d})}. \tag{41}$$

Also, let $\eta_0 = 0$ and $\eta_{\bar{d}+1} = 1$. We create $f = (f(0), f(1), ..., f(\bar{d}))$ so that

$$f(i) = \eta_{i+1} - \eta_i, \qquad \forall i = 0, 1, ..., \bar{d}. \tag{42}$$

Under each f out of the K possibilities, we test policies a for $T = 10,000$ periods on $L = 100$ sample paths of demand levels. In each period t on the lth path, we generate a random variable $\xi_{\bar{d}+(l-1)T+t}$ uniformly distributed on $[0, 1]$. Then, let demand d_t in that period be the only $d = 0, 1, ..., \bar{d}$ that satisfies

$$\eta_{d-1} \equiv \sum_{i=0}^{d-1} f(i) \leq \xi_{\bar{d}+(l-1)T+t} < \sum_{i=0}^{d} f(i) \equiv \eta_d. \tag{43}$$

Let y_t^a be the order-up-to level in period t under policy a. We use the following as an approximation to the policy's expected regret by time t:

$$r_{f,\beta}^{a,t} = \text{AVG}\{\sum_{s=1}^{t}[h \cdot (y_s^a - d_s)^+ + b \cdot (d_s - y_s^a)^+ - h \cdot (y_{f,\beta}^* - d_s)^+ - b \cdot (d_s - y_{f,\beta}^*)^+]\}, \tag{44}$$

where AVG stands for average over the $L = 100$ demand paths.

For any $\alpha \in (0,1)$, we let $R^{a,t}_{\alpha,\beta}$ be the conditional value at risk at the α-quantile of the K regrets $r^{a,t}_{f,\beta}$. For instance, $R^{a,100}_{95\%,\beta}$ would stand for the average of the top 50 highest $r^{a,100}_{f,\beta}$ values, where each is the regret of policy a under the designated β value by time $t = 100$, out of the $K = 1,000$ randomly generated f's. Also, $R^{a,t}_{0\%,\beta}$ would be the average of all the $r^{a,100}_{f,\beta}$'s of f's sampled from all over Δ. In addition, $R^{a,t}_{99.9\%,\beta}$ would be the worst average regret found over the $K = 1,000$ randomly generated distributions f. Suppose policy a is the $\mathbf{y}$ studied earlier. Then, when K and L both approach $+\infty$ and α approaches 100%, the value $R^{a,t}_{\alpha,\beta}$ will approach $\sup_{f\in\Delta} R^t_f(\mathbf{y})$, the focal point of Theorem 2, at the particular β. Since L is finite, each $r^{a,t}_{f,\beta}$ is merely an approximation of the regret $R^t_f(\mathbf{y})$. Moreover, the finiteness of K means that the f in Δ generating the worst regret will most likely be missed. Nevertheless, the $R^{a,t}_{\alpha,\beta}$ values will yield insights into regrets stemming from policies a.

We mainly test two policies a, with $a = 0$ and 1, respectively. Policy 0 is our newsvendor-based one defined through (12) and (13). Policy 1 is adapted from the SA-based one proposed by Huh and Rusmevichientong [11] to the current nonperishable-inventory setting. It still follows (12), although its generation of the $\hat{y}_t$'s is different from (13). For the latter, let $\varepsilon_t = \bar{d}/((h\vee b)\cdot\sqrt{t})$. There is also an auxiliary process $\hat{z}_t$. In the beginning, $y_1 = \hat{y}_1 = \hat{z}_1 = 0$. Then for $t = 2, 3, ...,$

$$\hat{y}_t = \begin{cases} \lfloor \hat{z}_t \rfloor, & \text{with probability } \lceil \hat{z}_t \rceil - \hat{z}_t, \\ \lceil \hat{z}_t \rceil, & \text{with probability } 1 - (\lceil \hat{z}_t \rceil - \hat{z}_t), \end{cases} \tag{45}$$

where

$$\hat{z}_t = \begin{cases} (\hat{z}_{t-1} - h\varepsilon_{t-1}) \vee 0 \wedge \bar{d}, & \text{when } \begin{array}{l} \hat{y}_{t-1} = \lfloor \hat{z}_{t-1} \rfloor \text{ and } d_{t-1} \le y_{t-1}, \text{ or} \\ \lfloor \hat{z}_{t-1} \rfloor < \hat{y}_{t-1} = \lceil \hat{z}_{t-1} \rceil \text{ and } d_{t-1} \le y_{t-1} - 1, \end{array} \\ (\hat{z}_{t-1} + b\varepsilon_{t-1}) \vee 0 \wedge \bar{d}, & \text{when } \begin{array}{l} \hat{y}_{t-1} = \lfloor \hat{z}_{t-1} \rfloor \text{ and } d_{t-1} \ge y_{t-1} + 1, \text{ or} \\ \lfloor \hat{z}_{t-1} \rfloor < \hat{y}_{t-1} = \lceil \hat{z}_{t-1} \rceil \text{ and } d_{t-1} \ge y_{t-1}. \end{array} \end{cases} \tag{46}$$

We have considered policy 2, which uses both (12) and the definition of the ε_t sequence. The policy is inspired by both stochastic approximation and Derman's [10] up-and-down idea. Instead of (13) or (45) and (46), it uses the following for its updating of $\hat{y}_t$:

$$\hat{y}_t = \begin{cases} \hat{y}_t - 1 \text{ with probability } h\varepsilon_{t-1}, & \text{when } d_{t-1} \le y_{t-1} - 1, \\ \hat{y}_{t-1} + 1 \text{ with probability } b\varepsilon_{t-1}, & \text{when } d_{t-1} \ge y_{t-1} + 1, \\ \hat{y}_{t-1} - \text{sgn}(h-b) \text{ with probability } |h-b|\cdot\varepsilon_{t-1}/2, & \text{when } d_{t-1} = y_{t-1}, \\ \hat{y}_{t-1}, & \text{otherwise.} \end{cases} \tag{47}$$

However, our simulation study indicates that policy 2 is not competitive against either of the previous two policies, except when β is close to 0.5, at which time it is better than policy 1 in some occasions. So we omit presenting its performances.

At various α and β values, and at different t points, we compare $R^{a,t}_{\alpha,\beta}$ among policies. At various β values, we can evaluate $R^{a,t}_{\alpha,\beta}$ for policies $a = 0$ and 1 at times $t = 1^2$, 2^2, ..., 100^2 and α values at 0%, 95%, and 99.9%. Except for scaling differences, we find the basic findings do not depend much on β. In Figures 1 to 3, we just present results at the three β values of 0.1, 0.5, and 0.9. Instead of t, we have used $\sqrt{t}$ as our horizontal axis.

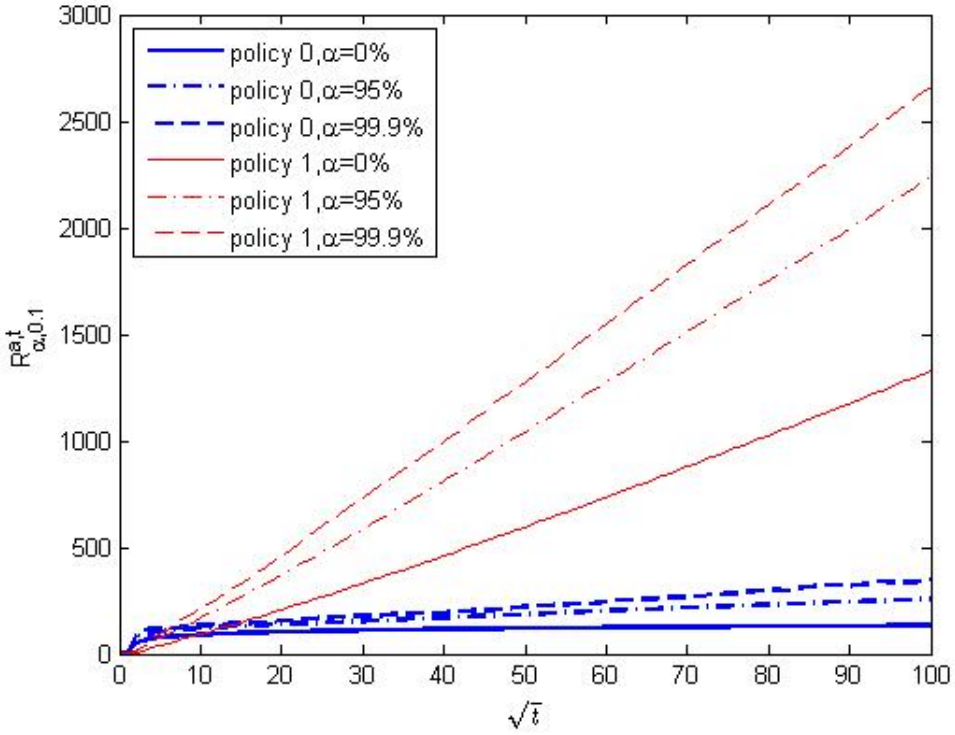

Figure 1: $R^{a,t}_{\alpha,\beta}$ Values when $\beta = 0.1$

From these figures, we can observe that policy 0 generates much smaller regrets, in both average and worst senses, than policy 1. This should be anticipated, as policy 0 utilizes more information regarding past observations than policy 1. We also see that regrets for both policies grow at approximately the rate of $\sqrt{t}$. In Figure 4, we provide a close-up for policy 0 at $\beta = 0.5$. Here, we have sampled $K = 10,000$ demand distributions, made measurements at time points $t = 1^2, 2^2, ..., 400^2$, and tried $\alpha = 0\%$, 95%, 99.9%, and 99.99%.

To understand what roles the distributions' separations from β have played in the for-

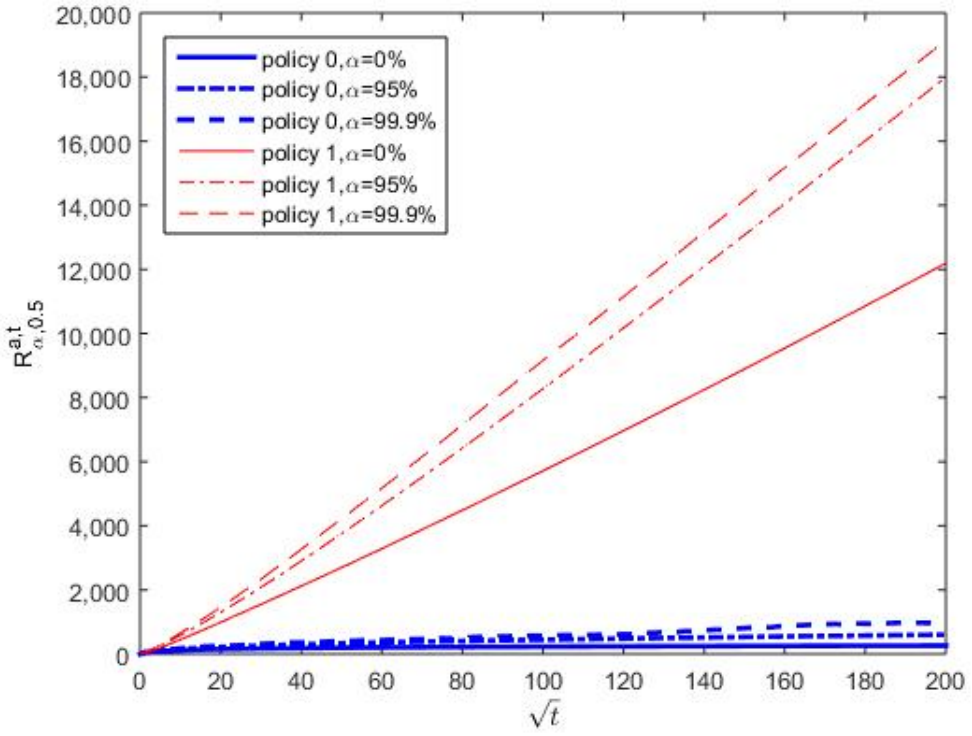

Figure 2: $R^{a,t}_{\alpha,\beta}$ Values when $\beta = 0.5$

mation of regrets, let us define $D^{a,t}_{\alpha,\beta}$ as the average of the separations $\delta_{f,\beta}$, as defined by (27) and (28), from among the $K \cdot (1-\alpha)$ distributions f with the worst regrets. For instance, $D^{a,100}_{95\%,\beta}$ would be the average of the $\delta_{f,\beta}$'s among the 500 f's which give the worst $r^{a,100}_{f,\beta}$ by time $t = 100$, out of the $K = 10,000$ randomly generated f's. In Figure 5, we present $D^{a,t}_{\alpha,\beta}$ for policies $a = 0$ and 1 at times $t = 1^2, 2^2, ..., 400^2$, α values at 0%, 95%, 99.9%, and 99.99%, and $\beta = 0.5$. Pictures at other β values look similar.

At $\alpha = 0\%$, the values $D^{a,t}_{0\%,0.5}$ are flat when t changes, because it is always the average of the $K = 10,000$ $\delta_{f,0.5}$'s. From Figure 5, we also see that $D^{0,t}_{\alpha,0.5}$ is decreasing in α at large enough t's. This is consistent with Theorem 1, showing that, in the long run, the selection of distributions f with large regrets will gravitate toward those with small separations from $\beta = 0.5$. In contrast, $D^{1,t}_{\alpha,0.5}$ seems to receive no clear influence from α.

We can also introduce an inseparability index $\gamma \in [0, 1)$ to indicate how difficult it is to separate demand distributions f from β. Suppose $d = 1, 2, ..., \bar{d}, \bar{d}+1$ with $\xi_{(d-1)} < \beta < \xi_{(d)}$ has been identified. Then, in the place of (41), we generate $\eta_1, ..., \eta_{d-1}$ so that

$$\eta_i = \frac{\xi_{(d-1)} + \gamma \cdot (\beta - \xi_{(d-1)})}{\xi_{(d-1)}} \cdot \xi_{(i)}, \qquad \forall i = 1, 2, ..., d-1, \tag{48}$$

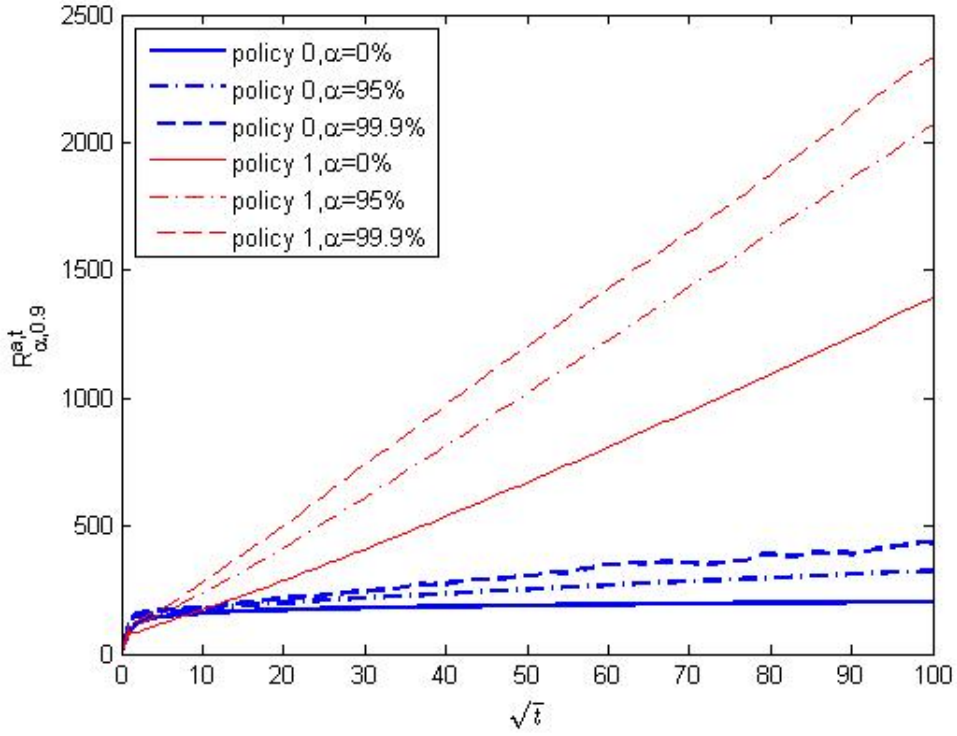

Figure 3: $R^{a,t}_{\alpha,\beta}$ Values when $\beta = 0.9$

and generate $\eta_d, ..., \eta_{\bar{d}}$ so that

$$\eta_i = 1 - \frac{1 - \xi_{(d)} + \gamma \cdot (\xi_{(d)} - \beta)}{1 - \xi_{(d)}} \cdot (1 - \xi_{(i)}), \qquad \forall i = d, d+1, ..., \bar{d}. \tag{49}$$

All other steps are the same as before. Our previous case happens when γ is set at the default value 0. When γ grows, there will be more chance that the generated f has a smaller separation from β. We can define $R^{a,t}_{\alpha,\beta,\gamma}$ as the α-quantile of the K regrets $r^{a,t}_{f,\beta}$, but this time with the f's generated under the given γ using (48) and (49) instead of (41). The previous $R^{a,t}_{\alpha,\beta}$ is just $R^{a,t}_{\alpha,\beta,0}$. In Figure 6, we show $R^{a,10000}_{\alpha,0.5,\gamma}$ for policies $a = 0$ and 1 at α values 0%, 95%, and 99.9%, and γ values at 0, 0.5, 0.9, 0.95, 0.99, 0.995, and 0.999.

Figure 6 shows that $R^{1,10000}_{\alpha,0.5,\gamma}$ is decreasing with γ. This is consistent with our earlier observation that smaller separations seem to help improve the performance of policy 1. At the same time, the increase of γ does not make policy 0 significantly worse. This seems to suggest a slow rise of the constant bound in Theorem 1 when the smallest allowed level of separation ϵ dwindles. For this to be reconcilable with the $\sqrt{t}$-sized rise in regret as indicated by both Theorem 2 and Figure 4, we will need the "leading" separation for policy 0 to decrease over time. This has somewhat been confirmed by Figure 5.

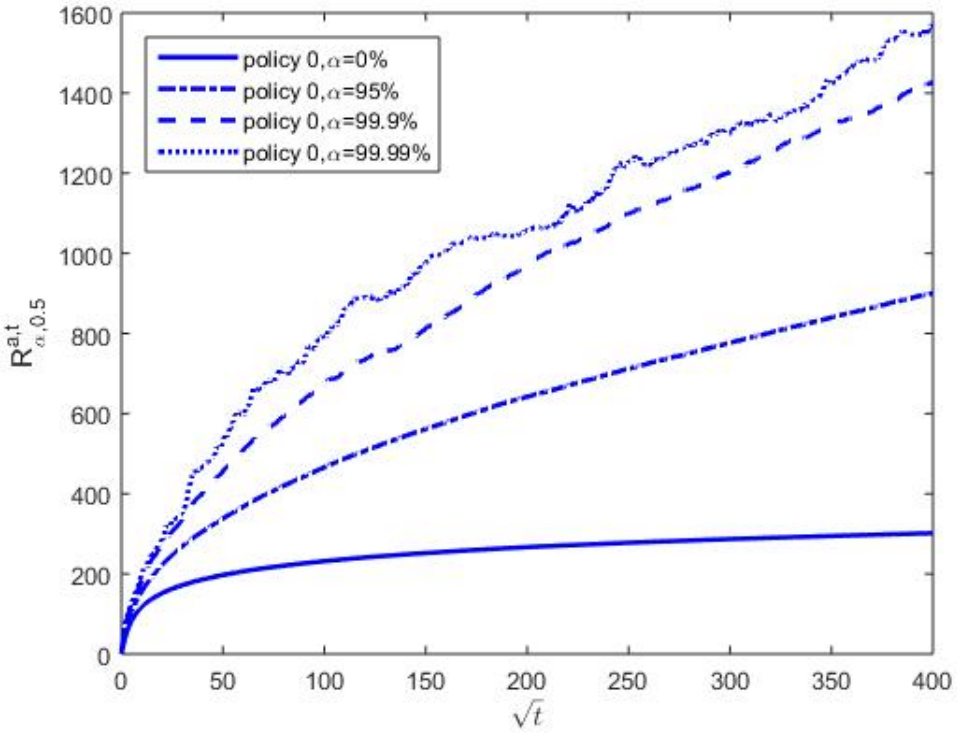

Figure 4: $R^{0,t}_{\alpha,\beta}$ Values when $\beta = 0.5$

6 Concluding Remarks

In regret bounds for a newsvendor-based adaptive policy, we have contributed to inventory control involving unknown demand of discrete nonperishable items. Currently, our analysis relies on knowledge of the maximum per-period demand $\bar{d}$. Also, our universal bound is related to a strictly positive parameter ϵ. Both deserve more attention. In addition, the newsvendor-based policy requires higher observability of historical demand levels than other policies, say the SA-based one. This makes it ill suited to situations involving demand censoring. Furthermore, we have not touched on realistic features like nonzero setup costs or lead times.

So a great deal awaits to be done in future research.

References

[1] Besbes, O. and A. Zeevi. 2009. Dynamic Pricing without Knowing the Demand Function: Risk Bounds and Near-optimal Algorithms. *Operations Research*, **57**, pp. 1407-

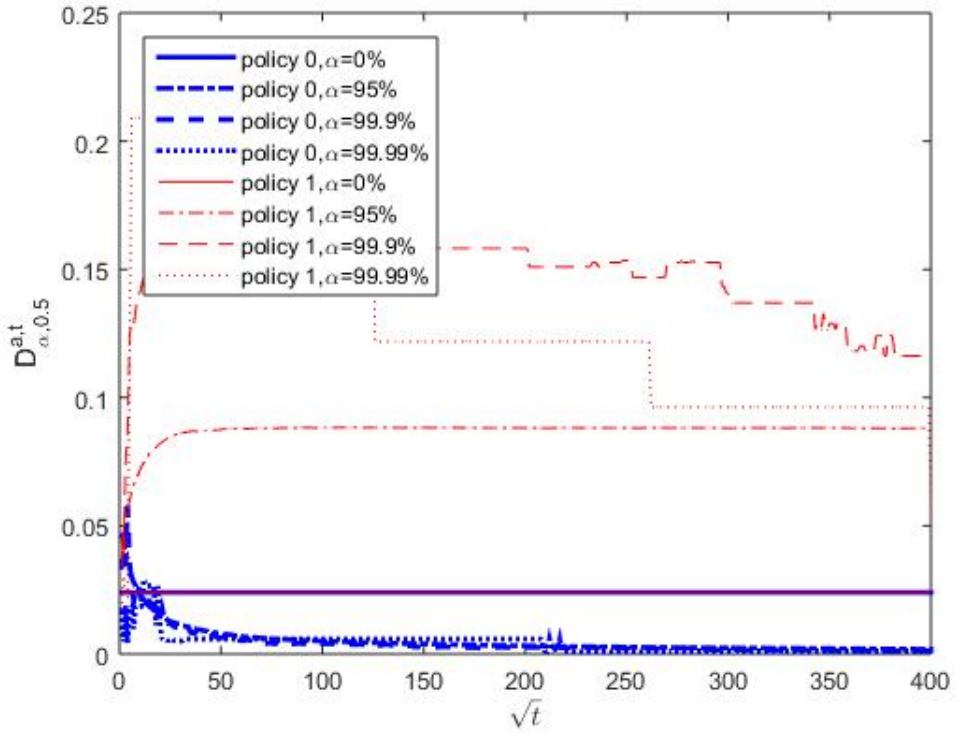

Figure 5: $D^{a,t}_{\alpha,\beta}$ Values when $\beta = 0.5$

1420.

[2] Besbes, O. and A. Muharremoglu. 2013. On Implications of Demand Censoring in the Newsvendor Problem. *Management Science*, **59**, pp. 1407-1424.

[3] Burnetas, A.N. and M.N. Katehakis. 1996. Optimal Adaptive Policies for Sequential Allocation Problems. *Advances in Applied Mathematics*, **17**, pp. 122-142.

[4] Burnetas, A.N. and M.N. Katehakis. 1997. Optimal Adaptive Policies for Markov Decision Processes. *Mathematics of Operations Research*, **22**, pp. 222-255.

[5] Burnetas, A.N. Kanavetas, O. and M.N. Katehakis. 2015. Asymptotically Optimal Multi-Armed Bandit Policies under a Cost Constraint, preprint arXiv:1509.02857.

[6] Cowan, W. and M.N. Katehakis. 2015. Asymptotically Optimal Sequential Experimentation Under Generalized Ranking, preprint arXiv:1510.02041.

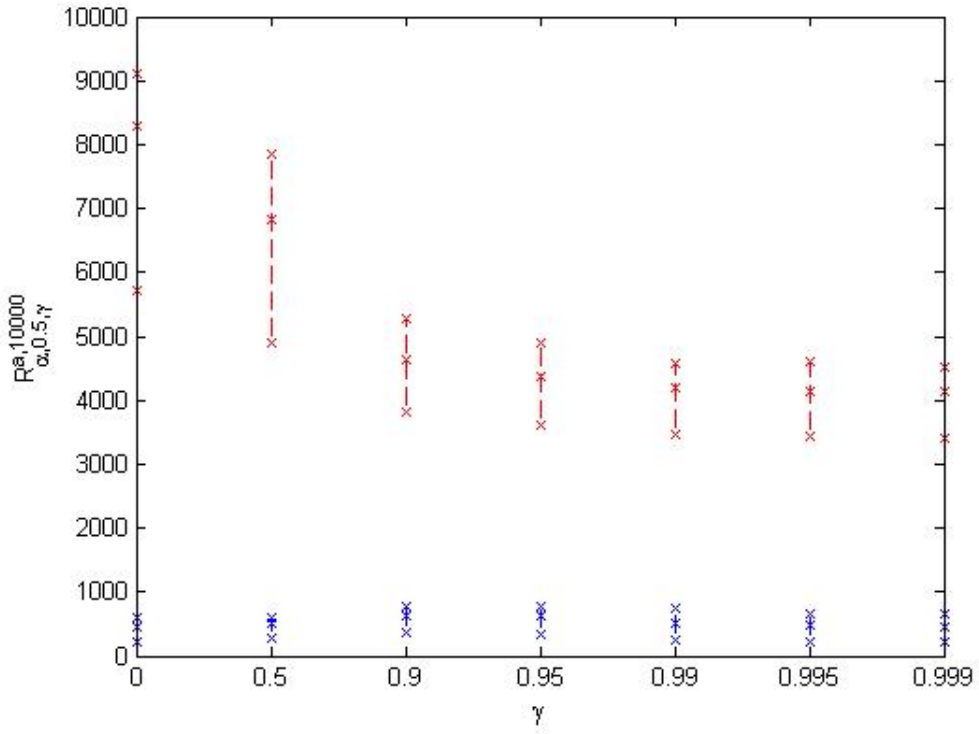

Figure 6: $R^{a,t}_{\alpha,\beta,\gamma}$ Values when $\beta = 0.5$ and $t = 10000$

[7] Cowan, W. and M.N. Katehakis. 2015. Asymptotic Behavior of Minimal-Exploration Allocation Policies: Almost Sure, Arbitrarily Slow Growing, Regret, preprint arXiv:1505.02865.

[8] Cover, T.M. and J.A. Thomas. 2006. *Elements of Information Theory, 2nd Edition*. Wiley-Interscience, New York.

[9] Dembo, A. and O. Zeitouni. 1998. *Large Deviations Techniques and Applications, 2nd Edition*. Springer, Heidelberg.

[10] Derman, C. 1957. Non-parametric Up-and-down Experimentation. *Annals of Mathematical Statistics*, **28**, pp. 795-798.

[11] Huh, W.T. and P. Rusmevichientong. 2009. A Non-parametric Asymptotic Analysis of Inventory Planning with Censored Demand. *Mathematics of Operations Research*, **34**, pp. 103-123.

[12] Katehakis, M.N. and H. Robbins. 1995. Sequential Choice from Several Populations. *Proceedings of the National Academy of Sciences*, **92**, pp. 8584-8585.

[13] Kiefer, J. and J. Wolfowitz. 1952. Stochastic Estimation of the Maximum of a Regression Function. *Annals of Mathematical Statistics*, **23**, pp. 462-466.

[14] Lai, T.L. and H. Robbins. 1985. Asymptotically Efficient Adaptive Allocation Rules. *Advances in Applied Mathematics*, **6**, pp. 4-22.

[15] Lariviere, M.A. and E.L. Porteus. 1999. Stalking Information: Bayesian Inventory Management with Unobserved Lost Sales. *Management Science*, **43**, pp. 346-363.

[16] Levi, R., R.O. Roundy, and D.B. Shmoys. 2007. Provably Near-optimal Sampling-based Policies for Stochastic Inventory Control Models. *Mathematics of Operations Research*, **32**, pp. 821-839.

[17] Perakis, G. and G. Roels. 2008. Regret in the Newsvendor Model with Partial Information. *Operations Research*, **56**, pp. 188-203.

[18] Robbins, H. 1952. Some Aspects of the Sequential Design of Experiments. *Bulletins of American Mathematical Society*, **58**, pp. 527-535.

[19] Robbins, H and S. Monro, S. 1951. A Stochastic Approximation Method. *Annals of Mathematical Statistics*, **22**, pp. 400-407.

[20] Scarf, H. 1959. Bayes Solutions of the Statistical Inventory Problem. *Annals of Mathematical Statistics*, **30**, pp. 490-508.

Appendices

A Proofs of Section 1

Proof of Proposition 1: Combining (5) and (15), we have

$$\begin{aligned} R_f^{T1}(\mathbf{y}) = h \cdot \textstyle\sum_{t=1}^{T} \{\mathbb{E}_f[(\hat{y}_t - D_t)^+] - \mathbb{E}_f[(y^*_{f,\beta} - D_t)^+]\} \\ + b \cdot \textstyle\sum_{t=1}^{T} \{\mathbb{E}_f[(D_t - \hat{y}_t)^+] - \mathbb{E}_f[(D_t - y^*_{f,\beta})^+]\}. \end{aligned} \tag{A.1}$$

When $x \geq y$,

$$(x-z)^+ - (y-z)^+ = (x - \max\{y,z\})^+; \tag{A.2}$$

whereas, when $x \leq y$,

$$(z-x)^+ - (z-y)^+ = (\min\{y,z\} - x)^+. \tag{A.3}$$

Plugging these into (A.1), we obtain

$$\begin{aligned} R_f^{T1}(\mathbf{y}) = h \cdot \textstyle\sum_{t=1}^{T} \sum_{y=y^*_{f,\beta}+1}^{\bar{d}} \mathbb{P}_f[\hat{y}_t = y] \cdot \mathbb{E}_f[(y - \max\{y^*_{f,\beta}, D_t\})^+] \\ + b \cdot \textstyle\sum_{t=1}^{T} \sum_{y=0}^{y^*_{f,\beta}-1} \mathbb{P}_f[\hat{y}_t = y] \cdot \mathbb{E}_f[(\min\{y^*_{f,\beta}, D_t\} - y)^+], \end{aligned} \tag{A.4}$$

which is below $(h\bar{d} + b\bar{d}) \cdot \sum_{t=1}^{T} \mathbb{P}_f[\hat{y}_t \geq y^*_{f,\beta} + 1] \vee \mathbb{P}_f[\hat{y}_t \leq y^*_{f,\beta} - 1]$. Therefore,

$$R_f^{T1}(\mathbf{y}) \leq (h\bar{d} + b\bar{d}) \cdot (\sum_{t=1}^{\tau_{f,\beta}} + \sum_{t=\tau_{f,\beta}+1}^{+\infty}) \mathbb{P}_f[\hat{y}_t \geq y^*_{f,\beta} + 1] \vee \mathbb{P}_f[\hat{y}_t \leq y^*_{f,\beta} - 1], \tag{A.5}$$

which, due to (22) to (24), is below the desired upper bound. ∎

Proof of Proposition 2: Combining (5) and (16), we have

$$R_f^{T2}(\mathbf{y}) = h \cdot \sum_{t=3}^{T} \mathbb{E}_f[(y_t - D_t)^+ - (\hat{y}_t - D_t)^+] + b \cdot \sum_{t=3}^{T} \mathbb{E}_f[(D_t - y_t)^+ - (D_t - \hat{y}_t)^+]. \tag{A.6}$$

Note $y_t \geq \hat{y}_t$ from (12). So in view of both (A.2) and the fact that $(d - y_t)^+ - (d - \hat{y}_t)^+ \leq 0$,

$$R_f^{T2}(\mathbf{y}) \leq h \cdot \sum_{t=3}^{T} \mathbb{E}_f[(y_t - \max\{\hat{y}_t, D_t\})^+]. \tag{A.7}$$

But this is further below $h \cdot \sum_{t=3}^{T} \mathbb{E}_f[(y_t - \hat{y}_t)^+]$. Hence,

$$R_f^{T2}(\mathbf{y}) \leq h\bar{d} \cdot \tau_{f,\beta} + h\bar{d} \cdot \sum_{t=\tau_{f,\beta}}^{+\infty} \mathbb{P}_f[y_{t+1} \geq \hat{y}_{t+1} + 1]. \tag{A.8}$$

So the key is to bound $\mathbb{P}_f[y_{t+1} \geq \hat{y}_{t+1} + 1]$ for $t \geq \tau_{f,\beta}$. However,

$$\begin{aligned}\mathbb{P}_f[y_{t+1} \geq \hat{y}_{t+1} + 1] &= \mathbb{P}_f[y_{t+1} \geq \hat{y}_{t+1} + 1 | \hat{y}_{t+1} = y^*_{f,\beta}] \cdot \mathbb{P}_f[\hat{y}_{t+1} = y^*_{f,\beta}] \\ &\quad + \mathbb{P}_f[y_{t+1} \geq \hat{y}_{t+1} + 1 | \hat{y}_{t+1} \neq y^*_{f,\beta}] \cdot \mathbb{P}_f[\hat{y}_{t+1} \neq y^*_{f,\beta}] \\ &\leq \mathbb{P}_f[y_{t+1} \geq y^*_{f,\beta} + 1] + \mathbb{P}_f[\hat{y}_{t+1} \leq y^*_{f,\beta} - 1] + \mathbb{P}_f[\hat{y}_{t+1} \geq y^*_{f,\beta} + 1].\end{aligned} \tag{A.9}$$

Thus, from (22) to (24),

$$\mathbb{P}_f[y_{t+1} \geq \hat{y}_{t+1} + 1] \leq \mathbb{P}_f[y_{t+1} \geq y^*_{f,\beta} + 1] + \exp(-\frac{\kappa_{f,\beta} \cdot (t - \tau_{f,\beta})}{2}). \tag{A.10}$$

Meanwhile,

$$\begin{aligned}\mathbb{P}_f[y_{t+1} \geq y^*_{f,\beta} + 1] &= \mathbb{P}_f[y_{t+1} \geq y^*_{f,\beta} + 1 | y_t \geq y^*_{f,\beta} + 1] \cdot \mathbb{P}_f[y_t \geq y^*_{f,\beta} + 1] \\ &\quad + \mathbb{P}_f[y_{t+1} \geq y^*_{f,\beta} + 1 | y_t \leq y^*_{f,\beta}] \cdot \mathbb{P}_f[y_t \leq y^*_{f,\beta}] \\ &\leq \mathbb{P}_f[y_t - D_t \geq y^*_{f,\beta} + 1 | y_t \geq y^*_{f,\beta} + 1] \cdot \mathbb{P}_f[y_t \geq y^*_{f,\beta} + 1] \\ &\quad + \mathbb{P}_f[\hat{y}_{t+1} \geq y^*_{f,\beta} + 1 | y_t \geq y^*_{f,\beta} + 1] \cdot \mathbb{P}_f[y_t \geq y^*_{f,\beta} + 1] \\ &\quad + \mathbb{P}_f[\hat{y}_{t+1} \geq y^*_{f,\beta} + 1 | y_t \leq y^*_{f,\beta}] \cdot \mathbb{P}_f[y_t \leq y^*_{f,\beta}] \\ &= \mathbb{P}_f[y_t - D_t \geq y^*_{f,\beta} + 1 | y_t \geq y^*_{f,\beta} + 1] \cdot \mathbb{P}_f[y_t \geq y^*_{f,\beta} + 1] + \mathbb{P}_f[\hat{y}_{t+1} \geq y^*_{f,\beta} + 1],\end{aligned} \tag{A.11}$$

where the first equality is an identity, the first inequality comes from (12), and the next equality is another identity when combining the previous two terms. Note $y_{t+1} = y_t - D_t$ or $y_{t+1} = \hat{y}_{t+1}$, and the latter would have occurred for sure when $y_{t+1} \geq y_t + 1$. Concerning part of the first term in the above, as $y_t \leq \bar{d}$,

$$\mathbb{P}_f[y_t - D_t \geq y^*_{f,\beta} + 1 | y_t \geq y^*_{f,\beta} + 1] \leq \mathbb{P}_f[D_t \leq \bar{d} - y^*_{f,\beta} - 1] = F_f(\bar{d} - y^*_{f,\beta} - 1). \tag{A.12}$$

Also in view of (22) to (24), we have from (A.11) and (A.12) that

$$\mathbb{P}_f[y_{t+1} \geq y^*_{f,\beta} + 1] \leq F_f(\bar{d} - y^*_{f,\beta} - 1) \cdot \mathbb{P}_f[y_t \geq y^*_{f,\beta} + 1] + \frac{\exp(-\kappa_{f,\beta} \cdot (t - \tau_{f,\beta})/2)}{2}. \tag{A.13}$$

Because $\epsilon_f \equiv f(\bar{d}) > 0$,

$$F_f(\bar{d} - y^*_{f,\beta} - 1) \leq F_f(\bar{d} - 1) = 1 - f(\bar{d}) = 1 - \epsilon_f < 1. \tag{A.14}$$

So for $t \geq \tau_{f,\beta}$, we have the iterative relation

$$\mathbb{P}_f[y_{t+1} \geq y^*_{f,\beta} + 1] \leq (1 - \epsilon_f) \cdot \mathbb{P}_f[y_t \geq y^*_{f,\beta} + 1] + \frac{\exp(-\kappa_{f,\beta} \cdot (t - \tau_{f,\beta})/2)}{2}. \tag{A.15}$$

Noting that $\mathbb{P}_f[y_{\tau_{f,\beta}} \geq y^*_{f,\beta} + 1] \leq 1$, we obtain

$$\mathbb{P}_f[y_{t+1} \geq y^*_{f,\beta} + 1] \leq (1 - \epsilon_f)^{t - \tau_{f,\beta} + 1} + \frac{1}{2} \cdot \sum_{s=0}^{t - \tau_{f,\beta}} (1 - \epsilon_f)^{t - \tau_{f,\beta} - s} \cdot \exp(-\frac{\kappa_{f,\beta} \cdot s}{2}). \tag{A.16}$$

Summing over $t = \tau_{f,\beta}, \tau_{f,\beta}+1, ...,$ we obtain

$$\sum_{t=\tau_{f,\beta}}^{+\infty} \mathbb{P}_f[y_{t+1} \geq y^*_{f,\beta}+1] \leq \frac{1-\epsilon_f}{\epsilon_f} + \frac{1}{2\epsilon_f \cdot (1-\exp(-\kappa_{f,\beta}/2))}. \tag{A.17}$$

Combining (A.8), (A.10), and (A.17), we obtain

$$\begin{aligned} R_f^{T2}(\mathbf{y}) &\leq h\bar{d} \cdot \tau_{f,\beta} + h\bar{d} \cdot \textstyle\sum_{t=\tau_{f,\beta}}^{+\infty} \mathbb{P}_f[y_{t+1} \geq y^*_{f,\beta}+1] \\ &\quad + h\bar{d} \cdot \textstyle\sum_{t=\tau_{f,\beta}}^{+\infty} \exp(-\kappa_{f,\beta} \cdot (t-\tau_{f,\beta})/2) \\ &\leq h\bar{d} \cdot \tau_{f,\beta} + h\bar{d} \cdot \textstyle\sum_{t=\tau_{f,\beta}}^{+\infty} \mathbb{P}_f[y_{t+1} \geq y^*_{f,\beta}+1] + h\bar{d}/(1-\exp(-\kappa_{f,\beta}/2)), \end{aligned} \tag{A.18}$$

which is below the desired upper bound. ∎

Proof of Theorem 1: Putting (14) as well as Propositions 1 and 2 together, we can obtain an upper bound for $R_f^T(\mathbf{y})$:

$$(2h\bar{d}+b\bar{d})\cdot\tau_{f,\beta} + \frac{3h\bar{d}+b\bar{d}}{2} \cdot \frac{1}{1-\exp(-\kappa_{f,\beta}/2)} + h\bar{d}\cdot\left(\frac{1-\epsilon_f}{\epsilon_f} + \frac{1}{2\epsilon_f \cdot (1-\exp(-\kappa_{f,\beta}/2))}\right). \tag{A.19}$$

For $f \in \Delta_{\epsilon,\delta,\beta}$, we already have $\epsilon_f \geq \epsilon$ by the definition in (29). Also, note that $\mathrm{D}_{KL}(u||\cdot)$ is convex with minimum achieved at u (Cover and Thomas [8], Theorems 2.6.3 and 2.7.2). So in view of (21), (27), and (28),

$$\kappa_{f,\beta} \equiv \min\{\mathrm{D}_{KL}(\beta||\alpha_{f.\beta}), \mathrm{D}_{KL}(\beta||\gamma_{f,\beta})\} \geq \min\{\mathrm{D}_{KL}(\beta||\beta-\delta_{f,\beta}), \mathrm{D}_{KL}(\beta||\beta+\delta_{f,\beta})\}, \tag{A.20}$$

which is further greater than $\kappa'_{\delta,\beta} \equiv \min\{\mathrm{D}_{KL}(\beta||\beta-\delta), \mathrm{D}_{KL}(\beta||\beta+\delta)\}$ due to f's membership in the $\Delta_{\epsilon,\delta,\beta}$ defined by (29). We can define $\tau'_{\delta,\beta}$ for $\kappa'_{\delta,\beta}$ in the same fashion in which $\tau_{f,\beta}$ is defined for $\kappa_{f,\beta}$; namely, through (23) and (24). Clearly, $\tau'_{\delta,\beta} \geq \tau_{f,\beta}$. Since (A.19) is decreasing in ϵ_f, $\kappa_{f,\beta}$, and increasing in $\tau_{f,\beta}$, we can replace them by, respectively, ϵ, $\kappa'_{\delta,\beta}$, and $\tau'_{\delta,\beta}$. The new right-hand side would constitute the desired $A_{\epsilon,\delta,\beta}$. ∎

B Proofs of Section 2

Proof of Proposition 3: For $f, g \in \Delta$, note that $Q_f(y^*_{g,\beta}) - Q_f(y^*_{f,\beta})$ can be written as

$$[Q_f(y^*_{g,\beta}) - Q_g(y^*_{g,\beta})] + [Q_g(y^*_{g,\beta}) - Q_g(y^*_{f,\beta})] + [Q_g(y^*_{f,\beta}) - Q_f(y^*_{f,\beta})]. \tag{B.1}$$

While the first and third terms can be made small when f and g are close, the second term is always negative due to $y^*_{g,\beta}$'s optimality when the underlying demand distribution is g. Let us investigate how small the first and third terms can be. For any $y = 0, 1, ..., \bar{d}$,

$$\begin{aligned} Q_f(y) - Q_g(y) &= h \cdot \textstyle\sum_{d=0}^{y-1}[f(d) - g(d)](y-d) + b \cdot \sum_{d=y+1}^{\bar{d}}[f(d) - g(d)](d-y) \\ &\leq (h \vee b) \cdot \bar{d} \cdot || \ f - g \ ||_1 = 2 \cdot (h \vee b) \cdot \bar{d} \cdot \delta_V(f, g). \end{aligned} \tag{B.2}$$

In view of the discussion around (B.1),

$$Q_f(y^*_{g,\beta}) - Q_f(y^*_{f,\beta}) \leq 4 \cdot (h \vee b) \cdot \bar{d} \cdot \delta_V(f, g). \tag{B.3}$$

Through Pinsker's inequality (33), we see that (B.3) will become

$$Q_f(y^*_{g,\beta}) - Q_f(y^*_{f,\beta}) \leq 2 \cdot (h \vee b) \cdot \bar{d} \cdot \sqrt{2\mathrm{D}_{KL}(g||f)}. \tag{B.4}$$

Note (5) also leads to

$$Q_f(y^*_{g,\beta}) - Q_f(y^*_{f,\beta}) \leq (h \vee b) \cdot \bar{d}. \tag{B.5}$$

For a fixed $\varepsilon > 0$, consider $G = \{g \in \Delta | \mathrm{D}_{KL}(g||f) \geq \varepsilon\}$. Then (31) will result with

$$\mathbb{P}_f[\mathrm{D}_{KL}(\hat{f}_{t-1}||f) \geq \varepsilon] \leq t^{\bar{d}+1} \cdot \exp(-\varepsilon \cdot (t-1)). \tag{B.6}$$

Consider $R^{T1}_f(\mathbf{y})$ defined at (15). By (13), we have

$$R^{T1}_f(\mathbf{y}) = \sum_{t=1}^{T} \{\mathbb{E}_f[Q_f(y^*_{\hat{f}_{t-1},\beta})] - Q_f(y^*_{f,\beta})\}. \tag{B.7}$$

Let ε_t be a sequence of positive constants. We then see that $R^{T1}_f(\mathbf{y})$ is below

$$\begin{aligned} &\textstyle\sum_{t=1}^{T} \{\mathbb{P}_f[\mathrm{D}_{KL}(\hat{f}_{t-1}||f) < \varepsilon_t] \cdot \mathbb{E}_f[Q_f(y^*_{\hat{f}_{t-1},\beta}) - Q_f(y^*_{f,\beta}) | \mathrm{D}_{KL}(\hat{f}_{t-1}||f) < \varepsilon_t] \\ &\qquad + \mathbb{P}_f[\mathrm{D}_{KL}(\hat{f}_{t-1}||f) \geq \varepsilon_t] \cdot \mathbb{E}_f[Q_f(y^*_{\hat{f}_{t-1},\beta}) - Q_f(y^*_{f,\beta}) | \mathrm{D}_{KL}(\hat{f}_{t-1}||f) \geq \varepsilon_t]\} \\ &\leq (h \vee b) \cdot \bar{d} \cdot \textstyle\sum_{t=1}^{T} [2\sqrt{2\varepsilon_t} + t^{\bar{d}+1} \cdot \exp(-\varepsilon_t \cdot (t-1))], \end{aligned} \tag{B.8}$$

where the inequality comes from (B.4) to (B.6).

Suppose $\varepsilon_t = (t-1)^{-1+2\epsilon}$ for some $\epsilon \in (0, 1/2)$. Plugging this into (B.8), we get

$$\begin{aligned} R^{T1}_f(\mathbf{y}) &\leq (h \vee b) \cdot \bar{d} \cdot \{2\sqrt{2} \cdot \textstyle\sum_{t=0}^{T-1} t^{-1/2+\epsilon} + \sum_{t=0}^{T-1} (t+1)^{\bar{d}+1} \cdot \exp(-t^{2\epsilon})\} \\ &= (h \vee b) \cdot \bar{d} \cdot \{2\sqrt{2} + 2\sqrt{2} \cdot T_1 + T_2\}, \end{aligned} \tag{B.9}$$

where

$$T_1 = \sum_{t=1}^{T-1} t^{-1/2+\epsilon}, \qquad \text{and} \qquad T_2 = \sum_{t=1}^{T-1} (t+1)^{\bar{d}+1} \cdot \exp(-t^{2\epsilon}). \tag{B.10}$$

Since $t^{-1/2+\epsilon}$ is decreasing in t,

$$T_1 \leq \int_0^{T-1} t^{-1/2+\epsilon} \cdot dt = \frac{2}{1+2\epsilon} \cdot (T-1)^{1/2+\epsilon}. \tag{B.11}$$

Note that $(t+1)^{\bar{d}+1} \cdot \exp(-t^{2\epsilon})$ is positive, increasing first, and decreasing next. So

$$T_2 \leq \int_0^{T-1} (t+1)^{\bar{d}+1} \cdot \exp(-t^{2\epsilon}) \cdot dt + \int_1^{T} (t+1)^{\bar{d}+1} \cdot \exp(-t^{2\epsilon}) \cdot dt. \tag{B.12}$$

Since $(t+1)^{2\epsilon} - t^{2\epsilon} < 1$, we have T_2 being below

$$\begin{array}{l} e \cdot [\int_0^{T-1} (t+1)^{\bar{d}+1} \cdot \exp(-(t+1)^{2\epsilon}) \cdot dt + \int_1^{T} (t+1)^{\bar{d}+1} \cdot \exp(-(t+1)^{2\epsilon}) \cdot dt] \\ \quad \leq 2e \cdot \int_1^{+\infty} t^{\bar{d}+1} \cdot \exp(-t^{2\epsilon}) \cdot dt, \end{array} \tag{B.13}$$

where e is the natural logarithmic base. Letting $u = t^{2\epsilon}$ and hence $t = u^{1/(2\epsilon)}$, we obtain

$$\int_1^{+\infty} t^{\bar{d}+1} \cdot \exp(-t^{2\epsilon}) \cdot dt = \frac{1}{2\epsilon} \cdot \int_1^{+\infty} u^{(\bar{d}+2-2\epsilon)/(2\epsilon)} \cdot \exp(-u) \cdot du. \tag{B.14}$$

Meanwhile, for $\alpha > 0$, integral by parts has $\int_1^{+\infty} u^{\alpha} \cdot \exp(-u) \cdot du$ equal to

$$e^{-1} + \alpha e^{-1} + \cdots + (\lfloor \alpha \rfloor - 1)! \cdot e^{-1} + (\lfloor \alpha \rfloor)! \cdot \int_1^{+\infty} u^{\alpha - \lfloor \alpha \rfloor - 1} \cdot \exp(-u) \cdot du, \tag{B.15}$$

which is below an α-dependent constant. Combining (B.13) to (B.15), we see that T_2 is below an ϵ-dependent constant. Together with (B.9) and (B.11), we see that for any $\epsilon \in (0, 1/2)$, there are constants A_ϵ and B_ϵ for the intended bound to be true. Furthermore, by adopting $A_\epsilon = A_{1/3}$ and $B_\epsilon = B_{1/3}$ for $\epsilon \geq 1/2$, the inequality can be maintained for every $\epsilon > 0$. ∎

Proof of Proposition 4: Let $\gamma = 1 - f(0)$. We divide the proof into two cases, with respectively, $\gamma \in [(1-\beta)/2, 1]$ and $\gamma \in [0, (1-\beta)/2)$.

Consider the first case with $\gamma \in [(1-\beta)/2, 1]$. For any positive integer τ_T, we show how $\mathbb{P}_f[S_{i+1} - S_i - 1 \geq \tau_T + 1]$ can be bounded. By the definition of the S_i's around (38),

$$\hat{y}_{S_{i+1}-1} \leq \hat{y}_{S_i} - D_{S_i} - D_{S_i+1} - \cdots - D_{S_{i+1}-2} - 1. \tag{B.16}$$

Since both $\hat{y}_{S_i}$ and $\hat{y}_{S_{i+1}-1}$ are between 0 and $\bar{d}$, the above necessitates that

$$D_{S_i} + D_{S_i+1} + \cdots + D_{S_{i+1}-2} \leq \bar{d} - 1. \tag{B.17}$$

This is only possible when there are at least $S_{i+1} - S_i - \bar{d}$ zeros among the $S_{i+1} - S_i - 1$ demand levels $D_{S_i}, D_{S_i+1}, ..., D_{S_{i+1}-2}$. When $S_{i+1} - S_i - 1 = \tau + 1 \geq \bar{d} - 1$, the latter event's

chance under f with $f(0) = 1-\gamma$ is, by the binomial formula,

$$\sum_{k=\tau-\bar{d}+2}^{\tau+1} \frac{(\tau+1)!}{k!\cdot(\tau+1-k)!}\cdot(1-\gamma)^k\cdot\gamma^{\tau+1-k} < (\tau+1)^{\bar{d}}\cdot(1-\gamma)^{\tau-\bar{d}+2}\cdot(1+\gamma+\cdots+\gamma^{\bar{d}}), \quad \text{(B.18)}$$

which is less than $(\tau+1)^{\bar{d}}\cdot(1-\gamma)^{\tau-\bar{d}+1}$. There exists $\theta_\gamma = \bar{d}-1, \bar{d}, ...$ so that when $\tau \geq \theta_\gamma$, the aforementioned term will decrease with τ. For $\tau_T \geq \theta_\gamma$, we can thus deduce that

$$\mathbb{P}_f[S_{i+1} - S_i - 1 \geq \tau_T + 1] < (\tau_T+1)^{\bar{d}}\cdot(1-\gamma)^{\tau_T-\bar{d}+1}. \quad \text{(B.19)}$$

Each of the terms in (39) is between 0 and $(h+b)\cdot\bar{d}$. So by (36) and (B.19),

$$\begin{aligned} R_f^{T2}(\mathbf{y}) \leq \textstyle\sum_{t=3}^{T}\sum_{s=2\vee(t-\tau_T)}^{t-1}\mathbb{E}_f[\mid Q_f(\hat{y}_s - D_s - \cdots - D_{t-1}) - Q_f(\hat{y}_t) \mid \times \\ \times\mathbf{1}(\hat{y}_t \leq \hat{y}_s - D_s - \cdots - D_{t-1} - 1)] \\ +(T-2)\cdot(h\bar{d}+b\bar{d})\cdot(\tau_T+1)^{\bar{d}}\cdot(1-\gamma)^{\tau_T-\bar{d}+1}. \end{aligned} \quad \text{(B.20)}$$

The above right-hand side can be written as

$$\sum_{\tau=1}^{\tau_T} R_f^{T2,\tau}(\mathbf{y}) + (T-2)\cdot(h\bar{d}+b\bar{d})\cdot(\tau_T+1)^{\,\bar{d}}\cdot(1-\gamma)^{\tau_T-\bar{d}+1}, \quad \text{(B.21)}$$

where for $\tau = 1, 2, ..., \tau_T$,

$$\begin{aligned} R_f^{T2,\tau}(\mathbf{y}) = \textstyle\sum_{t=\tau+2}^{T}\mathbb{E}_f[\mid Q_f(\hat{y}_{t-\tau} - D_{t-\tau} - \cdots - D_{t-1}) - Q_f(\hat{y}_t) \mid \times \\ \times\mathbf{1}(\hat{y}_t \leq \hat{y}_{t-\tau} - D_{t-\tau} - \cdots - D_{t-1} - 1)]. \end{aligned} \quad \text{(B.22)}$$

By (10) and (13), we have $\hat{y}_t \leq \hat{y}_{t-\tau} - D_{t-\tau} - \cdots - D_{t-1} - 1$ only if

$$\hat{F}_{t-\tau-1}(\hat{y}_{t-\tau}-1) < \beta \leq \hat{F}_{t-1}(\hat{y}_t) \leq \hat{F}_{t-1}(\hat{y}_{t-\tau} - D_{t-\tau} - \cdots - D_{t-1} - 1). \quad \text{(B.23)}$$

Also, due to the nature of the empirical distribution as illustrated in (11),

$$\hat{F}_{t-1}(\hat{y}_{t-\tau} - D_{t-\tau} - \cdots - D_{t-1} - 1) \leq \hat{F}_{t-1}(\hat{y}_{t-\tau}-1) \leq \hat{F}_{t-\tau-1}(\hat{y}_{t-\tau}-1) + \frac{\tau}{t-\tau}. \quad \text{(B.24)}$$

Therefore, $\hat{y}_t \leq \hat{y}_{t-\tau} - D_{t-\tau} - \cdots - D_{t-1} - 1$ only if

$$\beta \leq \hat{F}_{t-1}(\hat{y}_t) \leq \hat{F}_{t-1}(\hat{y}_{t-\tau} - D_{t-\tau} - \cdots - D_{t-1} - 1) < \beta + \frac{\tau}{t-\tau} \leq \beta + \frac{\tau_T}{t-\tau}, \quad \text{(B.25)}$$

an inequality alluded to earlier in (40). On the other hand, (5) has that, for $y \leq z$,

$$\begin{aligned} &Q_f(z) - Q_f(y) \\ &= h\cdot\textstyle\sum_{d=0}^{z-1}F_f(d) + b\cdot\sum_{d=z}^{\bar{d}-1}(1-F_f(d)) - h\cdot\sum_{d=0}^{y-1}F_f(d) - b\cdot\sum_{d=y}^{\bar{d}-1}(1-F_f(d)) \\ &= h\cdot\textstyle\sum_{d=y}^{z-1}F_f(d) - b\cdot\sum_{d=y}^{z-1}(1-F_f(d)) = (h+b)\cdot\sum_{d=y}^{z-1}(F_f(d)-\beta). \end{aligned} \quad \text{(B.26)}$$

Now by (B.22), (B.25), and (B.26), $R_f^{T2,\tau}(\mathbf{y})$ is less than $(h+b)\cdot\sum_{t=\tau+2}^{T}\mathbb{E}_f[Z_t]$, where

$$\begin{aligned} Z_t = &\textstyle\sum_{d=\hat{y}_t}^{\hat{y}_{t-\tau}-D_{t-\tau}-\cdots-D_{t-1}-1} \mid F_f(d)-\beta \mid \times \\ &\times\mathbf{1}(\beta\le \hat{F}_{t-1}(\hat{y}_t)\le \hat{F}_{t-1}(\hat{y}_{t-\tau}-D_{t-\tau}-\cdots-D_{t-1}-1)<\beta+\tau_T/(t-\tau)). \end{aligned} \tag{B.27}$$

In general,

$$Z_t \le \bar{d}. \tag{B.28}$$

Noting in turn that $\mid F_f(d)-\hat{F}_{t-1}(d)\mid \le \delta_V(f,\hat{f}_{t-1})$, which is less than $\sqrt{\mathrm{D}_{KL}(\hat{f}_{t-1}||f)/2}$ due to Pinsker's inequality (33), we also have

$$\begin{aligned} Z_t &\le \textstyle\sum_{d=\hat{y}_t}^{\hat{y}_{t-\tau}-D_{t-\tau}-\cdots-D_{t-1}-1}(\mid F_f(d)-\hat{F}_{t-1}(d)\mid + \mid \hat{F}_{t-1}(d)-\beta\mid)\times \\ &\quad\times\mathbf{1}(\beta\le \hat{F}_{t-1}(\hat{y}_t)\le \hat{F}_{t-1}(\hat{y}_{t-\tau}-D_{t-\tau}-\cdots-D_{t-1}-1)<\beta+\tau_T/(t-\tau)) \\ &\le \bar{d}\cdot[\sqrt{\mathrm{D}_{KL}(\hat{f}_{t-1}||f)/2}+\tau_T/(t-\tau)]. \end{aligned} \tag{B.29}$$

So for a sequence ε_t,

$$\begin{aligned} R_f^{T2,\tau}(\mathbf{y}) \le\ & (h+b)\cdot\textstyle\sum_{t=\tau+2}^{T}\mathbb{E}_f[Z_t|\mathrm{D}_{KL}(\hat{f}_{t-1}||f)\le\varepsilon_t]\cdot\mathbb{P}_f[\mathrm{D}_{KL}(\hat{f}_{t-1}||f)\le\varepsilon_t] \\ &+(h+b)\cdot\textstyle\sum_{t=\tau+2}^{T}\mathbb{E}_f[Z_t|\mathrm{D}_{KL}(\hat{f}_{t-1}||f)>\varepsilon_t]\cdot\mathbb{P}_f[\mathrm{D}_K(\hat{f}_{t-1}||f)>\varepsilon_t], \end{aligned} \tag{B.30}$$

which by (B.6), (B.28), and (B.29), is less than

$$(h\bar{d}+b\bar{d})\cdot\sum_{t=1}^{T}[\sqrt{\frac{\varepsilon_t}{2}}+t^{\bar{d}+1}\cdot\exp(-\varepsilon_t\cdot(t-1))+\frac{\tau_T}{t}]. \tag{B.31}$$

The situation we face is very similar to (B.8) except for the τ_T/t-term. So as in Proposition 3, for any $\epsilon>0$, there are constants C'_ϵ, D'_ϵ, and E so that

$$R_f^{T2,\tau}(\mathbf{y}) \le C'_\epsilon + D'_\epsilon\cdot T^{1/2+\epsilon/2} + E\tau_T\cdot\ln T. \tag{B.32}$$

Note the E-term stems from the τ_T/t-term in (B.31). In view of (B.20) and (B.21),

$$R_f^{T2}(\mathbf{y}) \le C'_\epsilon\tau_T + D'_\epsilon\tau_T\cdot T^{1/2+\epsilon/2} + E\tau_T^2\cdot\ln T + (T-2)\cdot(h\bar{d}+b\bar{d})\cdot(\tau_T+1)^{\bar{d}}\cdot(1-\gamma)^{\tau_T-\bar{d}+1}, \tag{B.33}$$

when τ_T is above the θ_γ defined right after (B.18). Otherwise, we have almost the same inequality, albeit with the last term replaced by $(T-2)\cdot(h\bar{d}+b\bar{d})$. Choose τ_T appropriately, say $\tau_T=\lfloor T^{\epsilon/2}\rfloor$. Then, as long as T is large enough, say greater than some $T^0_{\epsilon,\gamma}$, we can ensure that $\lfloor T^{\epsilon/2}\rfloor$ is above θ_γ. Very importantly, just because $\gamma\in(0,1]$, we can make sure that the last term, regardless whether $\lfloor T^{\epsilon/2}\rfloor$ is below or above θ_γ, is always bounded from above by a positive constant $F_{\epsilon,\gamma}$. Thus,

$$R_f^{T2}(\mathbf{y}) \le C'_\epsilon\cdot T^{\epsilon/2} + D'_\epsilon\cdot T^{1/2+\epsilon} + E\cdot T^\epsilon\cdot\ln T + F_{\epsilon,\gamma}. \tag{B.34}$$

However, as long as T is large enough, the $T^{1/2+\epsilon}$-sized term will dominate all other terms. A constant term can certainly cover the case when T is not that large. Therefore, positive constants $C_{\epsilon,\gamma}$ and D_ϵ exist for the intended inequality

$$R_f^{T2}(\mathbf{y}) \leq C_{\epsilon,\gamma} + D_\epsilon \cdot T^{1/2+\epsilon}. \tag{B.35}$$

Since $C_{\epsilon,(1-\beta)/2}$ can be used for cases with $\gamma \geq (1-\beta)/2 > 0$, we can have the intended bound, namely,

$$R_f^{T2}(\mathbf{y}) \leq C_\epsilon + D_\epsilon \cdot T^{1/2+\epsilon}, \tag{B.36}$$

as long as γ stays above $(1-\beta)/2$.

We now turn to the second case with $\gamma \in [0, (1-\beta)/2)$. From (39), $R_f^{T2}(\mathbf{y})$ is equal to

$$\begin{array}{c} \sum_{t=3}^{T} \sum_{s=2}^{t-1} \mathbb{E}_f[Q_f(\hat{y}_s - D_s - \cdots - D_{t-1}) - Q_f(\hat{y}_t) | L(t) = s] \cdot \mathbb{P}_f[L(t) = s] \\ \leq (h\bar{d} + b\bar{d}) \cdot \sum_{t=3}^{T} \max_{s=2}^{t-1} \mathbb{P}_f[\hat{y}_s - D_s - \cdots - D_{t-1} \geq 1], \end{array} \tag{B.37}$$

in which the inequality is attributable to the facts that $\sum_{s=2}^{t-1} \mathbb{P}_f[L(t) = s] \leq 1$ and that $Q_f(z) - Q_f(y) \leq (h\bar{d} + b\bar{d}) \cdot \mathbf{1}(z \geq y+1)$ for $z \geq y$, which is easy to see from (5). But

$$\mathbb{P}_f[\hat{y}_s - D_s - \cdots - D_{t-1} \geq 1] \leq \mathbb{P}_f[\hat{y}_s \geq 1] \wedge (\sum_{d=1}^{\bar{d}} \mathbb{P}_f[\hat{y}_s \geq d] \cdot \mathbb{P}_f[D_s + \cdots + D_{t-1} \leq d-1]). \tag{B.38}$$

Meanwhile, by (13) and the current range of γ,

$$\mathbb{P}_f[\hat{y}_s \geq 1] = \mathbb{P}_f[\hat{F}_{s-1}(0) < \beta] \leq \mathbb{P}_f[\delta_V(f, \hat{f}_{s-1}) > 1 - \beta - \gamma], \tag{B.39}$$

which, due to Pinsker's inequality (33), is below $\mathbb{P}_f[\mathrm{D}_K(\hat{f}_{s-1} || f) > 2 \cdot (1-\beta-\gamma)^2]$. But by Sanov's (31), the latter is below $s^{\bar{d}+1} \cdot \exp(-2 \cdot (1-\beta-\gamma)^2 \cdot (s-1))$. Thus,

$$\mathbb{P}_f[\hat{y}_s \geq 1] \leq s^{\bar{d}+1} \cdot \exp(-2 \cdot (1-\beta-\gamma)^2 \cdot (s-1)). \tag{B.40}$$

Now for $d = 1, 2, ..., \bar{d}$, let $\gamma_d = 1 - F_f(d-1)$. Our setup is such that $0 \leq \gamma_{\bar{d}} \leq \gamma_{\bar{d}-1} \leq \cdots \leq \gamma_1 = \gamma < (1-\beta)/2$. Again due to (13),

$$\mathbb{P}_f[\hat{y}_s \geq d] = \mathbb{P}_f[\hat{F}_{s-1}(d-1) < \beta] = \mathbb{P}_f[\sum_{\tau=1}^{s-1} \mathbf{1}(D_\tau \geq d) > (1-\beta) \cdot (s-1)]. \tag{B.41}$$

Note that $\mathbf{1}(D_1 \geq d), \mathbf{1}(D_2 \geq d), ..., \mathbf{1}(D_{s-1} \geq d)$ are independent Bernoulli random variables with mean γ_d, and hence $\sum_{\tau=1}^{s-1} \mathbf{1}(D_\tau \geq d)$ is a Binomial random variable with mean $\gamma_d \cdot (s-1)$. So by Markov's inequality, the rightmost term in (B.41) is below

$$\frac{\mathbb{E}_f[\sum_{\tau=1}^{s-1} \mathbf{1}(D_\tau \geq d)]}{(1-\beta) \cdot (s-1)} = \frac{\gamma_d}{1-\beta}. \tag{B.42}$$

Therefore,

$$\mathbb{P}_f[\hat{y}_s \geq d] \leq \frac{\gamma_d}{1-\beta}. \tag{B.43}$$

Also, it is easy to see that

$$\mathbb{P}_f[D_s + \cdots + D_{t-1} \leq d-1] \leq (1-\gamma_d)^{t-s}. \tag{B.44}$$

Combining (B.38), (B.40), (B.43), and (B.44), we obtain

$$\begin{aligned} &\max_{s=2}^{t-1} \mathbb{P}_f[\hat{y}_s - D_s - \cdots - D_{t-1} \geq 1] \\ &\quad \leq \max_{s=2}^{t-1}[s^{\bar{d}+1} \cdot \exp(-2\cdot(1-\beta-\gamma)^2 \cdot (s-1))] \wedge \\ &\qquad \wedge [\textstyle\sum_{d=1}^{\bar{d}} (\gamma_d/(1-\beta)) \cdot (1-\gamma_d)^{t-s}]. \end{aligned} \tag{B.45}$$

Consider $a(\gamma, \tau) \equiv \tau^{\bar{d}+1} \cdot \exp(-2\cdot(1-\beta-\gamma)^2 \cdot (\tau-1))$. There exists a $t_0 \geq 1$ so that for any $t \geq t_0$, the function $a((1-\beta)/2, t)$ will decrease with t, and also

$$a(\frac{1-\beta}{2}, t) = t^{\bar{d}+1} \cdot \exp(-\frac{(1-\beta)^2}{2} \cdot (t-1)) < \exp(-\frac{(1-\beta)^2 \cdot t}{4}). \tag{B.46}$$

Note also that $a(\gamma, s) < a((1-\beta)/2, s)$ for $\gamma \in (0, (1-\beta)/2)$. Next, consider $b(\gamma', \tau) \equiv \gamma' \cdot (1-\gamma')^\tau$. Note that

$$\frac{\partial b(\gamma', \tau)}{\partial \gamma'} = (1-\gamma')^{\tau-1} \cdot [1 - (\tau+1)\cdot\gamma'], \tag{B.47}$$

and

$$\frac{\partial^2 b(\gamma', \tau)}{\partial (\gamma')^2} = -\tau \cdot (1-\gamma')^{\tau-2} \cdot [2 - (\tau+1)\cdot\gamma']. \tag{B.48}$$

So the b-maximizing γ' is $\gamma^*_\tau = 1/(\tau+1)$. Plugging back, we have

$$b(\gamma^*_\tau, \tau) = \frac{1}{\tau+1} \cdot (1 - \frac{1}{\tau+1})^\tau = \frac{1}{\tau+1} \cdot \frac{1}{(1+1/\tau)^\tau}. \tag{B.49}$$

Note that $\lim_{\tau \to +\infty} (1+1/\tau)^\tau = e$, the natural logarithmic base which is above 2. So when τ is large enough, say greater than some t_1, the above will be below $1/(2\tau+2)$.

For $T \geq 2t_0 + 2t_1$, the upper bound in (B.45) is further bounded by a constant plus

$$\sum_{t=2t_0+2t_1+1}^{T} (\max_{s=2}^{\lfloor t/2 \rfloor} \sum_{d=1}^{\bar{d}} \frac{\gamma_d}{1-\beta} \cdot (1-\gamma_d)^{t-s} + \max_{s=\lfloor t/2 \rfloor + 1}^{t-1} s^{\bar{d}+1} \cdot \exp(-2\cdot(1-\beta-\gamma)^2 \cdot (s-1)), \tag{B.50}$$

which, according to the above from (B.46) to (B.49), is below

$$\sum_{t=2t_0+2t_1+1}^{T} [\frac{\bar{d}}{1-\beta} \cdot \max_{s=2}^{\lfloor t/2 \rfloor} \frac{1}{2t-2s+2} + \max_{s=\lfloor t/2 \rfloor + 1}^{t-1} \exp(-\frac{(1-\beta)^2 \cdot s}{4})]. \tag{B.51}$$

But this is smaller than

$$\sum_{t=2t_0+2t_1+1}^{T} [\frac{\bar{d}}{(1-\beta)\cdot t} + \exp(-\frac{(1-\beta)^2 \cdot t}{8})], \tag{B.52}$$

which has a constant-plus-$\ln T$ bound. So, there exist positive constants E and F so that

$$R_f^{T2}(\mathbf{y}) \leq E + F \cdot \ln T, \tag{B.53}$$

for any $\gamma \in (0, (1-\beta)/2)$. Now between (B.36) and (B.53), only the former has to be used when C_ϵ is made large enough. We therefore have the intended bound. ∎